CANADIAN THESES ON MICROFICHE

THÈSES CANADIENNES SUR MICROFICHE

National Library of Canada
Collections Development Branch

Canadian Theses on
Microfiche Service

Ottawa, Canada
K1A 0N4

Bibliothèque nationale du Canada
Direction du développement des collections

Service des thèses canadiennes
sur microfiche

NOTICE

The quality of this microfiche is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Previously copyrighted materials (journal articles, published tests, etc.) are not filmed.

Reproduction in full or in part of this film is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30. Please read the authorization forms which accompany this thesis.

AVIS

La qualité de cette microfiche dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

Les documents qui font déjà l'objet d'un droit d'auteur (articles de revue, examens publiés, etc.) ne sont pas microfilmés.

La reproduction, même partielle, de ce microfilm est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30. Veuillez prendre connaissance des formules d'autorisation qui accompagnent cette thèse.

THIS DISSERTATION
HAS BEEN MICROFILMED
EXACTLY AS RECEIVED

LA THÈSE A ÉTÉ
MICROFILMÉE TELLE QUE
NOUS L'AVONS REÇUE

NL 339 (r. 86/09)

Canada

National Library of Canada | Bibliothèque nationale du Canada

Ottawa, Canada
K1A 0N4

(95)

TC – 67348

ISBN 0-315-19940-5

CANADIAN THESES ON MICROFICHE SERVICE – SERVICE DES THÈSES CANADIENNES SUR MICROFICHE

PERMISION TO MICROFILM – AUTORISATION DE MICROFILMER

• Please print or type – Écrire en lettres moulées ou dactylographier

AUTHOR – AUTEUR

Full Name of Author – Nom complet de l'auteur

Ronald Harvey SMYTH

Date of Birth – Date de naissance	Canadian Citizen – Citoyen canadien
19/01/51	[X] Yes / Oui [ ] No / Non
Country of Birth – Lieu de naissance	**Permanent Address – Résidence fixe**
Canada	109 Gore Street #4 Kingston Ontario K7L 2L6

THESIS – THÈSE

Title of Thesis – Titre de la thèse

Cognitive Aspects of Anaphora Judgment and Resolution

Degree for which thesis was presented Grade pour lequel cette thèse fut présentée	Year this degree conferred Année d'obtention de ce grade
Ph.D.	1985
University – Université	**Name of Supervisor – Nom du directeur de thèse**
University of Alberta	Bruce L. Derwing

AUTHORIZATION – AUTORISATION

Permission is hereby granted to the NATIONAL LIBRARY OF CANADA to **microfilm** this thesis and to **lend** or **sell copies of the film.**

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

L'autorisation est, par la présente, accordée à la BIBLIOTHÈQUE NATIONALE DU CANADA de **microfilmer** cette thèse et de **prêter** ou de **vendre des exemplaires du film.**

L'auteur se réserve les autres droits de publication; ni la thèse ni de longs extraits de celle-ci ne doivent être imprimés ou autrement reproduits sans l'autorisation écrite de l'auteur.

ATTACH FORM TO THESIS – VEUILLEZ JOINDRE CE FORMULAIRE À LA THÈSE

Signature	Date
[signature]	April 22, 1985

NL-91 (r. 84/03)

Canada

THE UNIVERSITY OF ALBERTA

COGNITIVE ASPECTS OF ANAPHORA JUDGMENT AND RESOLUTION

by

© RONALD H. SMYTH

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE

DOCTOR OF PHILOSOPHY

IN

PSYCHOLINGUISTICS

DEPARTMENT OF LINGUISTICS

EDMONTON, ALBERTA

SPRING, 1985

THE UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR RONALD.HARVEY.SMYTH.....

TITLE OF THESIS COGNITIVE.ASPECTS.OF.ANAPHORA.JUDGMENT.AND.RESOLUTION.

DEGREE FOR WHICH THESIS WAS PRESENTED DOCTOR.OF.PHILOSOPHY.....

YEAR THIS DEGREE GRANTED 1985.....

Permission is hereby granted to THE UNIVERSITY OF ALBERTA LIBRARY to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

(Signed) .....

PERMANENT ADDRESS:

4-109.Gore.Street.....

Kingston,.Ontario.....

K7L.2L6.....

Canada.....

DATED October 25..... 1984

THE UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research, for acceptance, a thesis entitled "Cognitive Aspects of Anaphora Judgment and Resolution", submitted by Ronald H. Smyth in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Psycholinguistics.

.....

Supervisor

.....

.....

.....

External Examiner

.....

External Examiner

Date..October.25,.1984.....

For Chloe, St. Debbie

ABSTRACT

English pronominal anaphora is examined from three perspectives: (i) the role of processing strategies in grammaticality judgment; (ii) models of on-line resolution; and (iii) post-resolution effects on priming and activation.

On the basis of an overview of formal approaches to anaphora, a detailed analysis is presented of non-syntactic factors influencing coreference judgments by nonlinguists. It is argued that while naive speakers apply comprehension strategies during the judgment task, linguists are more likely to apply production frequency criteria. In a judgment study involving 33 examples of blocked coreference drawn from the literature, speakers allowed coreference on 55% of the items. Discussion focuses on methodological aspects of data collection for anaphora studies.

The experimental literature on anaphora resolution is also reviewed in detail. The exhaustive serial search model is challenged, and a new model incorporating on-line exploitation of syntactic and semantic cues is offered. This model implies that the "gender effect" (Springston, 1975; Kaplan, 1981) is an artifact of the use of verification tasks to study the resolution process, and an experiment separating reading times from verification times shows that the effect is attributable entirely to the latter task. A model of anaphora verification further predicts that the effect should appear only in

free anaphora, and is supported.

Finally, post-resolution activation effects are studied in a probe recognition task. It is shown that differences in activation for an antecedent and a word in the same proposition are due to selective partial deactivation of the "companion" word. This is interpreted as evidence for propositional representation and proposition-based retrieval even when surface codes are still available shortly after sentence processing. It is argued that ordinary reading tasks are appropriate for probe reaction studies, even with subject-controlled single sentence displays followed by a subject-controlled probe onset.

ACKNOWLEDGEMENTS

This thesis is an attempt to answer the question, "How do pronouns work?". The fact that I felt it both necessary and possible to answer it in two ways shows that my mentors at the University of Alberta have provided the kind of guidance I was seeking when I first learned of their commitment to the reunification of linguistics and psychology. The fact that the linguistic arguments are phrased in terms of cognitive processes, and that the psychological arguments stem in part from a linguistic analysis, is evidence of their influence.

I am especially indebted to Professors Bruce Derwing and Gary Prideaux, who taught me to explore the psychological implications of linguistic theories. They showed me that language is both simpler and more complicated than other linguists tend to think.

Of my many friends in Edmonton and Calgary, Tracey Derwing stands out as having made the greatest personal contribution to the completion of this work. The importance of her constant support, encouragement and prodding is inestimable. Her dog, however, I can live without.

Two other marvelous duos, Peter Assmann and Albi Calman, and Wendy Rollins and Graham Peacock, also extended friendship, advice, meals, and, in the latter case, even a home for the summer.

I would also like to thank Dr. John Darbyshire and Mrs. Jenny White of Hotel Dieu Hospital, Kingston, for allowing me to use hospital

facilities and staff in the running of my experiments. Also at Hotel Dieu, Luis Zegarra is the man most responsible for the completion of Experiments 2 and 3. Without the countless hours he spent trying to meet my changing programming needs on unfamiliar equipment, I could not have undertaken this type of research.

Sylvia Hains of the Department of Psychology at Queen's University put in long hours over this past summer running my statistical analyses and at the same time guiding me through yet another new operating system. Her patience and sense of humour were greatly appreciated. I am also grateful to that Department for providing the computing facilities and office space, and to my students for commenting on earlier versions of my ideas.

Finally, I would like to thank my other committee members for their valuable comments: Dr. Matthew Dryer of the Department of Linguistics, Dr. Peter Dixon of the Department of Psychology, and Professor Jim McCawley of the University of Chicago.

TABLE OF CONTENTS

CHAPTER		PAGE
I.	OVERVIEW	
1.1	Introduction	1
1.2	An Example	3
II.	FORMAL ACCOUNTS OF CONSTRAINTS ON ANAPHORA	5
2.1	Introduction	5
2.2	Definitional preliminaries	7
2.3	Introduction to syntactic accounts	10
2.4	Reflexives	10
2.5	Simple pronouns	15
2.6	Pronominalization and Reflexivization as Transformations	16
2.7	Interpretive theories of anaphora	19
2.8	Precede-and-command	20
2.9	Lasnik's Kommand analysis	23
2.10	Reinhart and c-command	26
2.11	Chomsky's binding conditions	30
2.12	C-command, adverbial scope, and process interpretations	32
2.13	Resolution biases and processing strategies	35
2.14	Anaphora and discourse structure	37
2.14.1	Anaphora without syntax	37
2.14.2	Kuno's discourse constraints	40

TABLE OF CONTENTS continued

		PAGE
2.15	Discourse cohesion effects: Relevance and topic shifts	45
2.16	Scene shifts	47
2.17	Discourse referents and discourse model synthesis	48
2.18	Summary of formal treatments of anaphora	50
III.	COREFERENCE JUDGMENTS, COMPREHENSION STRATEGIES, AND BLOCKED ANAPHORA	51
3.1	Introduction	51
3.2	Variables affecting coreference judgments	53
3.2.1	Assumed stress	53
3.2.2	Length	54
3.2.3	Perceived dialect variation	57
3.2.4	Connectives	58
3.2.5	Pragmatic content	59
3.2.6	Ungrammaticality due to factors other than coreference	61
3.2.7	Analogy and depth of processing	65
3.2.8	Judgments on isolated sentences: Task demands	67
3.2.9	Acceptability, preferred readings, and implicit context: more methodological pitfalls	68
3.3	Conclusions: Judgment, comprehension processes, and language acquisition	74

TABLE OF CONTENTS continued

	PAGE
IV. JUDGMENT EXPERIMENT	79
Introduction	79
Method	80
Subjects	80
Materials	80
Procedure	82
Results	82
Order of Test and Implication	82
Error rates	83
Means for YES, NO, and TEST items	83
Mean responses without regard to confidence	84
Mean confidence ratings	85
Reversals	86
Item analysis	90
Discussion	91
Summary of response patterns	91
Individual items	93
General discussion	96
V. PSYCHOLINGUISTIC APPROACHES TO ANAPHORA RESOLUTION	103
5.1 Introduction	103

TABLE OF CONTENTS continued

		PAGE
5.2	Experimental approaches to anaphora resolution	105
5.2.1	Pronouns vs. lexical NP's: Effects of definiteness	105
5.2.2	Sentence-level interpretive biases: Verb-based effects	109
5.2.3	Intersentential interpretive biases	117
5.2.4	Search processes	122
5.2.5	Complexity and dominance	125
5.2.6	Search domains	127
5.3	Reading times, verification times, and anaphora resolution	133
5.4	Anaphora resolution and the top-down bottom-up controversy	134
VI.	ON-LINE VS. POST-PROCESSING EFFECTS IN ANAPHORA RESOLUTION	141
6.1	Introduction	141
6.2	The Target reading model	142
6.2.1	Free forward Targets	143
6.2.2	Free backward Targets	144
6.2.3	Blocked forward Targets	145
6.2.4	Blocked backward Targets	145
6.3	The verification model	146
6.4	Experiment 2	149
	Method	149
	Stimulus passages and design	149
	Subjects	149
	Procedure	149

TABLE OF CONTENTS continued

	PAGE
Results	153
Pre-analysis	153
Target reading times	153
Verification times	154
Discussion	156
VII. POST-RESOLUTION PROCESSES: THE ACTIVATION OF ANTECEDENTS	161
7.1 Introduction: Priming and activation	161
7.2 Priming and activation in anaphora resolution	164
7.3 Activation of antecedents across clause boundaries: some residual problems	167
7.4 Experiment 3	169
Method	169
Materials and design	169
Subjects	170
Procedure	170
Results	171
Discussion	171
VIII. SUMMARY AND CONCLUSIONS	173
BIBLIOGRAPHY	176
APPENDIX A: Items, Experiment 1, List 1	188
APPENDIX B: Experiment 1: Instructions to subjects	193

TABLE OF CONTENTS continued

		PAGE
APPENDIX C:	Experiment 1, group 1: Response sheet	195
APPENDIX D:	Master List, Experiment 2	198
APPENDIX E:	SPART Listing, Experiment 2	207
APPENDIX F:	Instructions, Experiment 2	210
APPENDIX G:	Master List, Experiment 3	211
APPENDIX H:	SPART Listing, Experiment 3	218
APPENDIX I:	Instructions, Experiment 3	221
APPENDIX J:	Response Coincidence Analysis, Experiment 1	222

LIST OF TABLES

Table		Page
1.	Mean item scores and error rates, Experiment 1: Combined judgment-confidence scores	84
2.	Newman-Keuls test of differences among means, combined list scores, Experiment 1.	84
3.	Mean judgment scores, Experiment 1.	85
4.	Newman-Keuls test of differences among means, judgment scores, Experiment 1.	
5.	Mean confidence scores, Experiment 1.	85
6.	Newman-Keuls test, confidence ratings, Experiment 1.	86.
7.	Types arranged in homogeneous subgroups according to mean number of Reversals by Duncan's multiple range test.	86
8.	Scores for Test items, Experiment 1.	
9.	T-tests for paired comparisons, Test and matched Free items.	90
10.	Results of one-way ANOVA's for 3-element sets of syntactically related items, Experiment 1.	92.
11.	Target reading times per condition in milliseconds, Experiment 2. Based on medians.	154
12.	Blocking by Direction interaction, Experiment 2.	154
13.	Verification times and error rates per condition, Experiment 2, based on medians.	155
14.	Match by Blocking interaction, Experiment 2, Verification, based on medians.	155
15.	Mean probe latency in milliseconds, Experiment 3, with per cent correct and standard deviations.	171

List of Figures

Figure		Page
1.	Subject groups, Experiment 1.	225
2.	Item clusters across subject groups, Experiment 1.	226
3.	Item clusters for Group 1 subjects.	227
4.	Item clusters for Group 2 subjects.	228
5.	Item clusters for Group 3 subjects.	229
6.	Item clusters for Group 4 subjects.	230

CHAPTER ONE

OVERVIEW

1.1 Introduction

The problem of how English speakers are able to compute pronoun-antecedent relations, part of the broader anaphora resolution question, is receiving a good deal of attention from researchers in natural language syntax, discourse analysis, cognitive psychology, and artificial intelligence alike. Such an expenditure of time and energy might seem extravagant to the layman, for whom the process is entirely automatic, but attempts to model anaphora resolution have shown that it involves a surprisingly large number of morphological, lexical, syntactic, semantic, pragmatic, discourse, and suprasegmental variables. Moreover, identifying the many linguistic influences on coreference is only half the battle; one must also account for the implementation of these cues during the antecedent search process, in a manner which is compatible with existing models of word recognition, parsing, semantic integration, and memory. Anaphora resolution therefore provides a formidable challenge for the development of workable human and machine models of language comprehension.

This study begins, in Chapter II, with an extensive review and

elaboration of formal linguistic treatments of pronominal anaphora, followed by a discussion, in Chapter III, of several processing strategies which appear to influence whether or not speakers will judge coreference to be possible for a particular sentence. These strategies suggest that coreference judgments will be highly contaminated by factors irrelevant to the goal of specifying a speaker's output constraints on anaphoric relations. This hypothesis is tested, and confirmed, in the judgment experiment reported in Chapter IV.

Chapter V reviews the information processing literature on anaphora resolution, examining the factors governing the many search, assignment, activation, and integration processes triggered when a pronoun is encountered in text. This includes such questions as the type of search involved (serial vs. parallel; exhaustive vs. self-terminating), when resolution begins (during processing of the pronoun itself, or later in the clause, sentence or discourse) and ends (again, during pronoun processing or later), the type of memory model required (semantic activation of antecedents), search domains (effects due to clause or sentence boundaries, foregrounding, and the number of potential preceding and following antecedents), search priorities (which potential antecedents are checked sooner or more quickly), search direction (forward vs. backward), and many other related issues. Chapters VI and VII present two further experiments concerning pronoun processing. Experiment 2 compares search processes for antecedents which both precede and follow the pronoun, and for both syntactically permitted and disallowed antecedents. The results of this study are interpreted from the perspective of competing top-down and bottom-up models of the comprehension process. Experiment 3

demonstrates that even in very short-term searches (through an immediately preceding clause), a pronoun activates both its antecedent and NP concepts in the same clause, a result which is not entirely accounted for by Kintsch's (1974) model of the propositional representation of text in memory.

The final chapter consists of a general discussion.

1.2 An example

Consideration of the following pair of sentences will assist the reader in gaining a preliminary grasp of the complexity of anaphora resolution.

1. Vietnam is currently seeking advice from other countries on agricultural matters.

2. Due to its enormous population, China has the kind of experience in food production that Vietnam needs.

How does a reader or listener determine that the intended antecedent of its is China? Resolution appears to depend on at least the following factors:

(i) its is singular and neuter (its ≠ countries): MORPHOLOGY

(ii) its may refer to an NP either within or outside the sentence in which it appears (its = either Vietnam or China): SYNTAX

(iii) Vietnam is the initial topic of discourse, and the only country mentioned when the pronoun is encountered (its = Vietnam is a preferred reading): DISCOURSE STRUCTURE.

(iv) neither experience nor food production may be said to have a population (its ≠ experience or food production): LEXICAL SEMANTICS

(v) the population of China is enormous, while the population of Vietnam is not (its ≠ Vietnam): WORLD KNOWLEDGE

These are but some of the cues to correct resolution to be discussed in the following chapters.

Such an analysis also raises a variety of interesting questions concerning processing. For example, if resolution is attempted on-line, (2) should be a garden-path sentence. Moreover, one would like to be able to specify at what point an initial hypothesis (its = Vietnam) begins to disintegrate: is processing disrupted immediately in response to incoming information concerning the population of the asserted antecedent, or are such complexities deferred to a later stage, such as the next clause boundary? Further questions concern the nature of the search: Are only NP's scanned? Is it the contents of working memory that are canvassed? Does distance from the pronoun have an effect? Are backward searches easier than forward searches? What is the time required to conduct the search?

This single example suggests that the study of anaphora resolution is not only a fruitful enterprise, but also one which gives us fascinating insights into the detailed structure of language comprehension.

CHAPTER TWO

FORMAL ACCOUNTS OF CONSTRAINTS ON ANAPHORA

2.1 Introduction

Anaphoric processes were given scant treatment by linguists prior to the 1960's. Early analyses (e.g., Bloomfield, 1933; Jespersen, 1909-1949; Hockett, 1955; Gleason, 1955) were concerned mainly with the morphological aspects of pronoun systems. With respect to the syntactic phenomena, they made only passing reference to the fact that pronouns "replace" or "stand for" full NP's (see Crymes, 1968). Examination of syntactic constraints was begun by Lees and Klima (1963), and certain of the resolution biases for otherwise free anaphora were recognized by Jackendoff (1973). However, as recently as 1973, Quirk, Svartvik and Greenbaum say of contextual factors in resolution only that "we interpret the appropriate substitutions from the content of the sentence" (p. 294).

Syntactic analyses attempt to characterize such restrictions in terms of one or two principles of interpretation from surface phrase structure; they differ mainly in their definitions of the

dominance relations on which these principles depend. Much of this work constitutes an attempt to clarify the role of surface phrase structure relations such as precedence and command in blocking coreference between personal pronouns and candidate antecedents (e.g., Chomsky, 1980, 1981; Langacker, 1966, 1969; Reinhart, 1976, 1981; Wasow, 1979). Wasow provides this simplified account:

> Given an NP and a definite pronoun in the same sentence, the NP may serve as the antecedent for the pronoun unless...the pronoun is to the left of the NP and is less deeply embedded than the NP (p. 61).

Coreference is thus blocked for he and John in (3) below, where he must to refer to some other, unspecified antecedent, and free in (4), where, given an appropriate context, he may refer either to John or to some other NP:

3. *He said that John was sick.

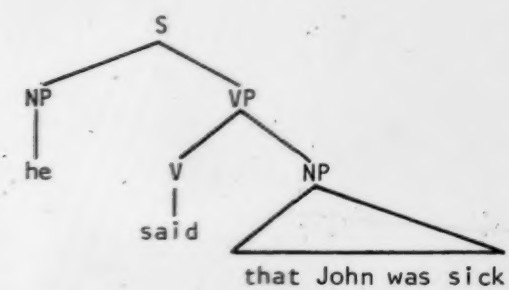

4. John said that he was sick.

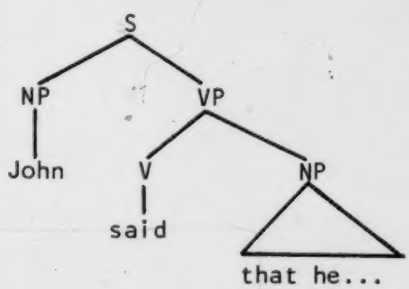

(Tree structures are simplified; asterisks indicate blocked coreference between underlined NP's)

According to the most popular classification scheme, the blocked coreference is said to run backward from John to he in (3), while (4) is termed free forward coreference.

Interpretive theories within the transformational paradigm date

from at least Jackendoff (1968), and in general have attempted to provide a comprehensive account of sentence-level syntactic constraints which block coreference between specified NP's. Because of the nature of the data on which the earlier accounts were based, this was for several years referred to as the Backward Anaphora Constraint. Later recognition of the existence of blocked forward anaphora has led to the demise of this term.

It should be emphasized that in syntactic accounts, lexical, semantic, pragmatic, discourse, suprasegmental, and processing variables are treated as nuisance variables modulating coreference interpretation within a set of fundamental syntactic constraints. This view is far from unanimous: more pessimistic accounts (e.g., Bily, 1977; Bolinger, 1977, 1979) view the entire syntactic approach as a red herring, and claim that counter-examples are so numerous as to cast serious doubt on the very assumption that anaphora is a syntactic process.

2.2 Definitional preliminaries

Perhaps a good deal of potential terminological confusion can be forestalled by examining here the various uses of the term anaphora by linguists.

Bloomfield (1933) introduces the terms anaphora and antecedent in a chapter entitled 'Substitution phenomena':

> ...the form for which substitution is made has occurred in recent speech. Thus, when we say Ask that policeman, and he will tell you, the substitute he means among other things that the singular male substantive expression which is replaced by he, has been recently uttered. A substitute which implies this, is an anaphoric or dependent substitute, and the recently-uttered replaced form is the antecedent. (p. 249)

Thus for Bloomfield, a pronoun must have an explicit NP antecedent, which precedes it in discourse (Bloomfield does not discuss backward anaphora), and which may or may not be in the same sentence as the pronoun.

Waterhouse (1963), on the other hand, defines anaphora as a purely inter-sentential phenomenon: an anaphoric pronoun is one "without identified referent in the same sentence" (p. 46), and until quite recently, transformationalists defined anaphora as a purely intra-sentential phenomenon. Dougherty (1969) states that a pronoun

> ...with anaphoric interpretation...can be understood as being coreferential with some other noun phrase in the sentence (p.)

In transformational terms, a sentence containing a free pronoun is often said to be ambiguous between the anaphoric and non-anaphoric readings (but cf. Lasnik, 1976), while a sentence with blocked coreference is said to have only a non-anaphoric reading. Thus, clauses containing reflexive and reciprocal pronouns normally have only an anphoric reading. This terminology, according to which anaphoricity is a property attributed to entire sentences, is problematic for cases such as

5. John told Fred to hit him.

in which the pronoun is ambiguous between the anaphoric and non-anaphoric readings with respect to John, but has only a non-anaphoric reading with respect to Fred. That is, anaphoricity is not a property of the sentence as a whole, but only of the syntactic relations between NP's.

Halliday and Hasan (1976) distinguish between endophoric and

exophoric reference, reserving the latter term for situational ("deictic" or "pragmatic") reference, and the former for reference within a text. Endophoric reference is further categorized according to whether the antecedent appears earlier (anaphoric) or later (cataphoric) in the text, but these terms do not distinguish between inter- and intra-sentential coreference.

Most recently, Chomsky (e.g., 1980, 1981) and others working within the Government and Binding framework have made a different distinction between anaphoric and pronominal expressions. The former include reflexives and reciprocals, for which coreference is syntactically determined, and the latter include the personal pronouns regardless of whether coreference is free or blocked in particular sentence in which they appear.

In the following pages, I shall take the term anaphora to mean a relation of either possible (in the case of simple pronouns) or necessary (for reflexives and reciprocals) coreference between two NP's, defined in terms of linguists', or native speakers' judgments and usage. Where convenient, I adopt Dougherty's distinction between anaphoric and non-anaphoric sentences, as well as Halliday and Hasan's finer distinctions.

A final point of definition concerns the nature of reference itself, whose explication is obviously beyond the scope of this work. Following Lyons (1977), I shall occasionally adopt what has become colloquial usage among linguists, according to which a pronoun is said to refer to its antecedent. Of course, this is not strictly correct, since the essential characteristic of (non-deictic) pronouns

is that they share a referent ("co-refer") with their antecedent. The antecedent itself, being a linguistic expression, is not the referent.

2.3 Introduction to syntactic accounts

As a first step toward constructing a model of anaphora resolution, we turn now to attempts at characterizing the linguistic environments in which coreference is possible, impossible, probable, or indeterminate. One source of such information is the literature dealing with syntactic constraints on pronominal coreference.

In examining this literature, two caveats are in order. First, none of the theories currently available is "correct": all make a certain number of wrong coreference predictions. This fact in itself bodes poorly for the popular notion that coreference options are governed by a single principle of interpretation from surface phrase structure. Secondly, there is much disagreement as to the data to be accounted for. While some of these disagreements are discussed openly in the literature, I will demonstrate that judgments of coreference options are exceedingly susceptible to a variety of influences, and further, that such indeterminacy is an unavoidable consequence of the fact that anaphora is at once a syntactic, semantic, and discourse phenomenon.

2.4 Reflexives

In this section I present several miscellaneous facts about coreference constraints on reflexives in order to determine the viability of the interpretive approach to describing them.

The simplest statement of this is the Clausemate Constraint, first proposed (as a transformational rule) by Lees and Klima (1963).

Lees and Klima noted that when two NP's in the same simplex sentence are coreferential, the rightmost NP is obligatorily a reflexive pronoun of appropriate gender, number and person:

6. Albi showered herself with gifts.

Conversely, a reflexive pronoun obligatorily has a clausemate as its antecedent. Thus, the ungrammaticality of (7):

7. *Peter showered herself with gifts.

is attributable not to a lack of antecedent for herself, but to the assertion of coreference between the gender-mismatched Peter and herself.

Lees and Klima claimed that the constraint holds without exception only when the reflexive functions as a direct or indirect object; in other positions the constraint is relaxed with certain verbs:

8. John ignored the oil on him.

9. John smeared the oil on himself.

However, this is an overly simplistic account, since the PP's of (8) and (9) differ in other ways. Since the PP of (8) can be viewed as a reduced relative clause, the use of a non-reflexive can be attributed to the fact that John and him are not in the same simplex sentence, as required by their statement of the rule.

Nonetheless, reflexives are, in most cases, in complementary distribution with simple pronouns, which cannot be coreferential with a clausemate:

10. *Peter saw him in a mirror.

11. *Patti bought her a new wading pool.

(Note that the asterisks indicate only that the coreferential readings are impossible, and not that the sentences are otherwise ungrammatical).

Ross (1970) notes a class of exceptions to the Clausemate Constraint: a reflexive pronoun may, in some cases, have its antecedent in the next highest clause if and only if there is no intervening clausemate NP which agrees with it in animacy. This is true of the so-called "picture noun" sentences such as

12. Tom believes that there is a picture of himself hanging in the post office.

In addition, coreference can skip over the picture noun only to the next highest NP which agrees in animacy:

13. *Tom believes that Sheila saw a picture of himself.

In a sense, then, a picture noun plus PP modifier behaves as though the head of the PP were a clausemate of the nearest higher non-picture NP.

Two additional facts are relevant to this account. First, consider the case of two potential antecedents in the higher clause, both matching the reflexive in animacy:

14. The waitress showed Sheila a picture of herself.

Jackendoff (1972) offers similar examples of the apparent tendency to interpret such sentences with resolution to the subject. That this is merely a bias is demonstrated by sentences such as

15. The waiter showed Sheila a picture of herself.

in which a gender mismatch with the subject allows object resolution. In addition, it is interesting to ask how a speaker would express (14)

with herself = Sheila: would one simply use the ambiguous form, hoping to clear up the vague reference later, or would one use the simple pronoun?

16. The waitress showed Sheila a picture of her.

Secondly, consider the case of an extra intervening inanimate noun: can coreference skip both?

17. Lorne noticed that the slide showed a sketch of him/himself.

I have encountered mixed reactions to such sentences. Most speakers say that both are possible, but they disagree as to which is better. Those who have a fairly strong preference for the simple pronoun (myself included) find the following cases perfectly acceptable:

18. Lorne told Jack that the slide showed a sketch of himself.

19. Lorne told Jack that the slide showed a sketch of him.

This minimally contrasting pair suggests that despite the split opinions over (17), coreference can indeed skip over two inanimate NP's, and that the reflexive/simple pronoun contrast is functional when only a resolution strategy can accomplish interpretation: reflexives bias the interpretation to the higher subject; non-reflexives bias interpretation to the non-subject. I do not believe that this has been discussed in the literature on reflexives.

The fact that speakers disagree on this point is in itself rather interesting, for it suggests that the rarity of picture noun sentences, and especially of these specially chosen examples, can result in non-application of the interpretive principles governing the simpler cases. The existence of speakers who reject the distinction shown in (18) - (19) raises the possibility that speakers are not necessarily able

to generalize interpretive principles to unusual cases. If this is so, the validity of collecting naive speakers' judgments on crucial cases is placed in jeopardy.

Chomsky (1981) shows that in certain subjectless complement clauses, the Clausemate Constraint is relaxed to allow coreference with an NP in a higher clause:

20. John appears to be enjoying himself.

His trace theory analysis, however, remains compatible with the constraint if one assumes that the reflexive pronoun is coreferential with the trace left by movement of John out of the lower clause:

21. John appears t to be enjoying himself.

Finally, Ross (1970) presents an example in which an animate NP may be the antecedent for a pronoun in a lower clause, even with intervening animate NP's:

22. Albert accused me of having tried to get his secretary to tell you that physicists like himself were hard to get along with.

Only if the intervening NP's match the reflexive pronoun on all relevant features does complete ambiguity result:

23. Albert accused Fred of having tried to get Mark to tell Michael that physicists like himself were hard to get along with.

Despite the possibility of coreference with any of the proper names, only Albert and Michael seem to be likely antecedents. This can be attributed to the pragmatics of the verb accuse: he = Fred is pragmatically bizarre, since criticizing oneself is not the sort of reprehensible act of which one might normally be "accused" (a similar argument holds for him = Mark). On the other hand, the reading is not

entirely excluded, since Fred could have some complex motive for wanting to convince someone else of his own character flaws - perhaps he wishes to get out of a prior agreement to work with Michael. It would appear that the amount of inferencing required in order to construct such a context somehow precludes the two marked readings. Of course, this too brings forth a number of questions that are relevant to both human and machine comprehension models: how much inferencing is actually done before a reading is adopted? Would the non-preferred reading still be more difficult to arrive at in a context in which it is plausible? What type of memory model could account for such context effects?

To summarize, then, the Clausemate Constraint on the interpretation of direct and indirect object reflexive pronouns provides a syntactically based, deterministic resolution strategy: the pronoun is obligatorily coreferential with its clausemate, subject to certain lexical and syntactic biases.

Since I have approached reflexives from the standpoint of interpretive principles, I should now reiterate that Lees and Klima originally formulated the constraint as a transformation which replaced a lexical NP by a reflexive pronoun. This will be clarified in the following sections.

2.5 Simple pronouns

The third person pronouns (he, she, it, they (nominative); his, her, its, their (genitive); him, her, it, them (accusative)) are not subject to deterministic resolution principles. For example,

although a simple pronoun cannot have a clausemate antecedent:

24. *John saw him in the mirror.

the antecedent is unspecified. Perceiving that coreference is blocked is only one step toward resolution; other structural principles governing the free/blocked distinction are the topic of the following sections.

2.6 Pronominalization and reflexivization as transformations

Lees and Klima (1963) proposed that pronouns are introduced by a transformational rule of Pronominalization (or Reflexivization) which replaces an underlying full NP with a pronoun under morphological identity. Chomsky (1965) found it necessary to alter this approach in order to conform with the Katz-Postal Hypothesis (Katz and Postal, 1964), which states that meaning is interpreted from deep structure; this accounts for the fact that pairs such as

25. John hurt John.

26. John hurt himself.

may differ in interpretation: (25) is ambiguous between the anaphoric and the non-anaphoric reading, but (26) has only an anaphoric reading. Since in the Lees-Klima formulation the deep structure of (26) is roughly (25), the theory would incorrectly predict that (26) is also ambiguous. Chomsky therefore proposed that deep structures be altered so as to allow the generation of referential indices along with underlying NP's. The two readings of (25) would then derive from distinct deep structures differing in whether or not the two occurrences of John are coindexed. By adding the condition that

Pronominalization and Reflexivization include a referential identity condition, one can derive (26) unambiguously from the deep structure in which the NP's are coindexed.

Over the next few years, however, it became apparent that the transformational analysis was untenable for other reasons. Some of the main arguments for the alternative interpretivist theory (in which coreference is determined from derived phrase marker configurations) will be presented here.

Perhaps the best-known (but no longer accepted; see Perlmutter and Soames, 1979) argument against the transformationalist position is the so-called Bach-Peters paradox (Bach, 1970), which shows that if the identity condition is to be maintained for co-indexing, the assumption that phrase markers must have a finite number of branches is violated. Consider (27):

27. The pilot who shot at it hit the mig that chased him.

Here, it would be coindexed with the mig that chased him, and him would be co-indexed with the pilot who shot at it. However, each of these full NP's contains a pronoun which must be co-indexed with another NP which contains a pronoun, and so on, ad infinitum. Under the assumption that infinite deep structures are disallowed, the transformational analysis must be rejected.

A second argument (Lakoff, 1968, 1969) concerns pronominal epithets - NP's which (when unstressed) may be coreferential with a lexical antecedent:

28. I told Ms. Barham about the sink, but the stupid bitch wouldn't do anything about it.

Since these NP's seem to obey all the distributional restrictions

that apply to pronouns, one would expect a theory of anaphora to handle them with the same formal mechanism. However, the transformational treatment allows coreference only under deep structure identity; the Pronominalization rule would have to be altered so as to allow for substitution of epithets rather than pronouns in such cases. This poses a problem in that epithets also carry meaning, which cannot be interpreted from the later stage in the derivation. In other words, a transformational approach to the interpretation of pronominal epithets violates the Katz-Postal hypothesis.

A third set of arguments (Bresnan, 1970) is based on the interaction of Pronominalization with movement rules such as Particle Movement and Dative Movement. Since these rules apply cyclically, it is possible for an NP to be moved within a lower clause before an antecedent in a higher clause is "seen" by the Pronominalization rule, and for that NP to be pronominalized after it has been moved into a position where pronouns are disallowed:

29. The girl claims that John turned the girl away (Deep structure).

30. The girl claims that John turned away the girl. (Particle Movement)

31. *The girl claims that John turned away her. (Pronominalization).

The only way to prevent this would be to impose an ad hoc constraint preventing Pronominalization from applying to structures to which certain other transformations have already applied.

Finally, Dougherty (1969) argues that in treating anaphoric pronouns differently from non-anaphoric pronouns, the transformational analysis fails to account for the fact that for every sentence

with an anaphoric reading, there is also a non-anaphoric reading (Wasow, 1979, notes several counter-examples). Thus, the two readings of

32. John said he was sick.

have very different derivational histories (one with an underlying non-coreferential pronoun, the other with a co-indexed John), and the fact that (32) is ambiguous appears to be accidental. On the other hand, a theory which interprets coreference from surface phrase structure and generates all pronouns in deep structure would simply label (32) as a case of free anaphora and correctly interpret it as ambiguous between the anaphoric and the non-anaphoric readings. Since Chomsky (1971) had also argued that interpretation of meaning from surface structure is necessary in the case of topic and focus, this innovation was ultimately adopted, and the Katz-Postal Hypothesis laid to rest.

Further arguments against the transformational treatment may be found in Wasow (1979), Jackendoff (1968), and elsewhere.

2.7 Interpretive theories of anaphora

None of the above arguments against the transformational model poses a problem for an interpretive theory of anaphora. In interpretive theories, all pronouns are introduced by lexical insertion and coreference is determined on the basis of derived structural configurations. This solves the problems inherent in the transformational analyses: the Bach-Peters sentences are generated with pronouns in deep structure, and the problem of infinite recursion does not arise; the meanings of pronominal epithets are available in deep structure, and coreference interpreted at a later stage; and independent con-

straints prevent the application of the offending movement rules to block sentences such as (31).

As has already been suggested, a finding that all cases of free anaphora share a structural property not found in blocked cases would be of great interest, particularly if that property could be shown to be reflected in aspects of linguistic performance, such as processing complexity, language acquisition, or judgment data. The next three sections examine some specific proposals concerning such interpretive principles.

2.8 Precede-and-command

Langacker (1966, 1969) represents an early attempt to formalize the constraints on sentence-internal coreference in terms of the horizontal and vertical positions of pronouns and NP's in phrase structure. His constraint is that:

> A pronoun may not both precede and command its antecedent. ...A node A "commands" another node B if (1) neither A nor B dominates the other; and (2) the S-node that most immediately dominates A also dominates B. (p. 167)

Langacker considered precedence and command to be the "primacy" relations relevant to anaphora, in that a pronoun which bears neither relation to an NP may be coreferential to (33), as may one which bears one of the relations (34,35). However, if both relations hold, coreference is blocked:

33. Ralph is more intelligent than he looks.

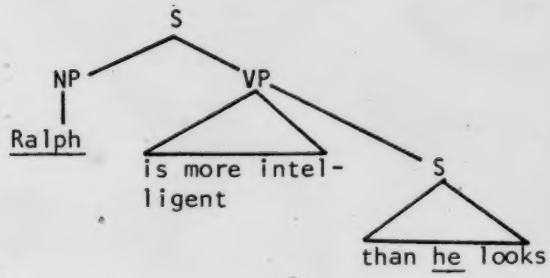

He neither precedes nor commands Ralph.

34. The woman who is to marry Ralph will visit him tomorrow.

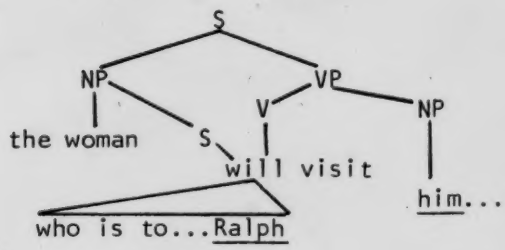

Him commands, but does not precede Ralph.

35. The woman who protected him hates Peter.

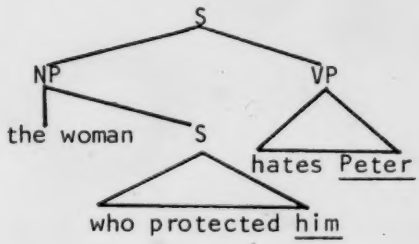

Him precedes, but does not command Peter.

36. *He hates the woman who rejected Peter.

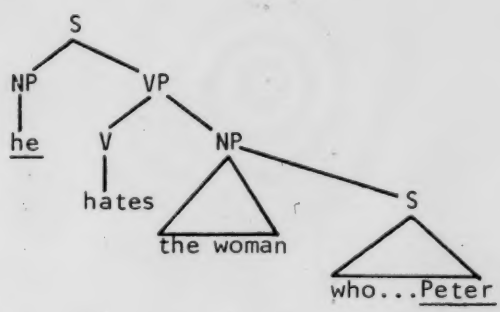

He both precedes and commands Peter.

While the constraint accounts for a broad range of data, a number of counter-examples demonstrate that it is not entirely correct. First, it cannot account for all instances of acceptable backward coreference:

37. Near him, Dan saw a snake. (Lakoff, 1968)

38. The chairman hit him on the head before the lecturer had a chance to say anything.

39. Irene sent letters to him while Fred was out of town.

40. I stayed with his wife while Fred was out of town.

In each case, the pronoun precedes and commands the antecedent, but the coreferential readings are possible.

Although Langacker failed to examine preposed adverbials such as (37), he implicitly accepts a prediction of ungrammaticality for (38) - (40) by starring (41):

41. * I knew him when Harvey was a little boy.

This judgment of noncoreference has been disputed by Reinhart (1981), and Carden (1978) cites several occurrences of such structures in spontaneous text. Thus the theory fails in these cases as well.

Precede-and-command also fails to block forward coreference of any kind:

42. *Near Dan, he saw a snake.

43. *How obnoxious to Ben's friends he is.

where the pronoun commands, but does not precede its antecedent.

It should also be noted that while the constraint blocks backward anaphora in some sentences which are otherwise handled by the Clausemate Constraint:

44. *He shot John.

it cannot account for forward blocking in similar structures:

45. *John shot him.

In (45) him does not precede John, yet coreference is blocked; the Clausemate Constraint is still necessary. Given the inadequacy of Precede-and-Command, one is left to speculate as to whether a single rule will emerge to replace it, or whether constraints on anaphora are

best handled by a larger number of rules.

2.9 Lasnik's Kommand analysis

Lasnik (1976) attempted to improve upon Langacker's analysis by providing an account of the acceptability of coreference in the case of genitive NP's which precede and command their antecedent (e.g., 40 above). This argument is based on Chomsky's (1969) observation that NP's can be cyclic nodes, in that they may head a domain for the application of transformations such as Passive:

46. The city was destroyed by the enemy.

47. The city's destruction by the enemy...

48. The enemy destroyed the city.

49. The enemy's destruction of the city...

Lasnik observed that the dominance of a genitive NP by a higher NP node in structures such as (50):

50.

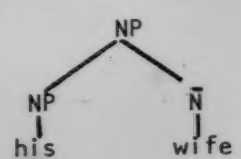

can be viewed as a barrier to command via the higher S node, as in (51):

51.

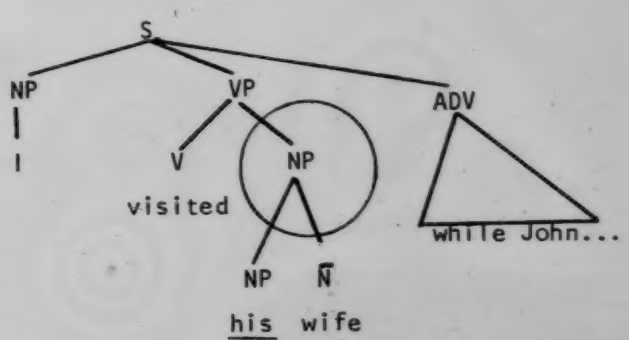

The circled NP node, being a cyclic node, dominates his but not

John, and no longer commands John under Lasnik's definition of the kommand relation:

52. A kommands B if the minimal cyclic node dominating A also dominates B. (p. 15)

Langacker's constraint is then altered as follows:

53. If NP_1 both precedes and Kommands NP_2, and NP_2 is not a pronoun, then coreference between NP_1 and NP_2 is blocked.

It is interesting to note that Lees and Klima (1963) saw a similar connection between cyclicity and pronominalization in their analysis of sentences such as (54):

54. Mary's father supported her / *herself.

which appear to be counter-examples to the Clausemate Constraint. They suggested that since "genitives in English are transformational in origin" (p. 22), the source of (54) is roughly (55):

55. The father$_{\#}$the father is Mary's$_{\#}$supported Mary.

with only one occurrence of Mary in the matrix sentence, and hence no reflexivization. Thus Lees and Klima realized that genitive NP's do not participate in coreference relations in the same way as other NP's in a main clause.

O'Grady (1983) deals with these exceptions in a partly semantic manner, claiming that genitives, as "peripheral" nominals with respect to the verb (subjects, objects and indirect objects are "core" nominals in his system), are permitted to be coreferential with the later NP because they do not violate his Precedence Constraint:

56. A core pronoun may not precede a peripheral antecedent.

It would appear that what is handled by dominance principles in the other theories is subsumed under the semantically-based distinction

between core and periphery.

We therefore have good evidence that genitives form a coherent class of exceptions to Langacker's constraint. However, Lasnik's analysis leaves a curious asymmetry in accounting for the grammaticality of (51), but not of similar structures such as (57):

57. I visited him while John was in Edmonton.

As we shall soon see, the grammaticality of structures such as (57) has been hotly debated. However, if (57) is grammatical even for only a subset of speakers, the cyclic nodes analysis cannot account for its existence. But this leads to two somewhat unsatisfactory conclusions: either the cyclic nodes analysis needs to be supplemented by yet another ad hoc specification of certain NP's which are not cyclic nodes, but which allow coreference in analogous structures, or the entire notion of cyclicity is not relevant to anaphora, and needs to be abandoned in favour of a set of exception rules to supplement precede-and-command. A third alternative, that command is simply the wrong relation, seems to be supported by the fact that the structure of (40) and of (57) share identical dominance relations despite the lack of cyclic nodes in the latter.

To carry the discussion one step further, note that the special NP-over-NP structure of Lasnik's cyclic nodes examples is also found in conjuncts; perhaps it is this abstract phrase marker configuration which affects the interpretation of coreference. For example, the following pair differ only in the presence or absence of the conjunct, yet (58) allows coreference, while (59) does not:

58. I told her and the producers that Meryl would be playing Sophie.

59. *I told her that Meryl would be playing Sophie.

(The argument becomes tenuous when we take into account the difficulty of collecting confident judgments on (58); see Chapter 4). On the other hand, conjuncts and genitives differ with respect to their interaction with the Clausemate Constraint: genitives, as we have already seen, are immune to that constraint, while conjuncts are not:

60. *John saw him in the mirror.

61. John saw his wife in the mirror.

62. *John saw him and Fred in the mirror.

Finally, even this argument is weakened somewhat by the fact that sentences such as (62) can be found in nonstandard dialects, and, I believe, in the language of children. This suggests that there may indeed be a link between conjuncts and genitives with respect to anaphora. The question of whether (40) and (57) allow coreference for the same reasons is thus unclear, but if that is indeed the case, the Kommand analysis cannot capture the generalization.

2.10 Reinhart and c-command

Reinhart (1976, 1981), while rejecting the precede-and-command analysis, nevertheless maintains that anaphora restrictions can be accounted for by a single surface interpretive principle. In addition to handling the blocked forward and free backward exceptions to precede-and-command, her constraint is also intended to account for certain asymmetries between coreference options for subjects and non-subjects, as in (63) and (64):

63. * In front of Dan, he saw a snake.

64. In front of Dan, I saw his snake.

This asymmetry poses a problem for any analysis incorporating command or Kommand, since everything c(k)ommanded by a subject is also c(k)ommanded by the object. However, simply restricting the application of some version of the Langacker analysis would be insufficient, since non-subjects can also be excluded from coreference, at least according to Reinhart's judgment on (65):

65. *I'm willing to give him two grand for Ben's car.

One very serious problem to be approached in these chapters is that of disagreement among linguists as well as among naive speakers as to which sentences in fact disallow coreference. Notice that Reinhart's analysis is largely justified by judgments such as (65); I will soon demonstrate that this and many other of her judgments are suspect, at least in terms of speaker's ability to judge the grammaticality of coreference in the crucial examples.

Reinhart's proposal is that

66. A given NP cannot be interpreted as coreferential with a distinct nonpronoun in its c-command domain. (p. 617)

where c-command is defined as follows:

67. Node A c(onstituent)-commands node B iff the branching node α_1 most immediately dominating A either dominates B or is immediately dominated by a node α_2 which dominates B, and α_2 is of the same category type as α_1. (p.612)

C-command is simply the converse of Klima's (1964) in construction with relation, which Culicover (1976) used to account for free backward pronominalization (backward pronominalization is permissible if the NP is not in construction with the pronoun). Reinhart's insight is

that c-command is equally applicable to forward pronominalization, as in (68) and (69):

68. * Near Dan, he saw a snake. 69. Near him, Dan saw a snake.

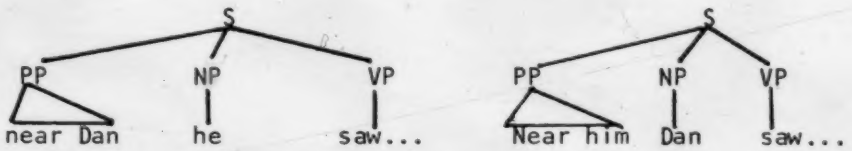

In (68), Dan does not c-command he, but he c-commands Dan; in (69) the opposite is true. In (68), Dan is in the domain of he (where domain of a node refers to all the nodes that it c-commands); coreference is blocked. In (69), however, he is in the domain of Dan, but is not a pronoun; coreference is allowed. Notice in particular that the c-command formulation makes no reference to precedence.

As further examples, consider again the tree structures of (35) and (36) (page 21). In (35), no c-command relation exists between the pronoun and its antecedent, and restriction (66) does not block coreference. In (36), he c-commands Peter, and coreference is blocked by (66).

Since c-command can be shown to subsume the command relation, the only points of divergence between Reinhart's analysis and those of Langacker or Lasnik involve sentences in which (a) A c-commands B and B precedes A (since forward coreference is blocked, an impossibility under the other analyses), and (b) A precedes and commands B, but does not c-command it (such that only c-command correctly allows coreference). The former has already been illustrated in (68), where A = he and B = Dan (such sentences with preposed PP's are discussed

in detail below); the latter involves sentences such as (70), which Langacker judged to be ungrammatical:

70. We fired him, since McIntosh's weird habits had finally reached an intolerable stage.

Notice that (70) has the same structure as (57), which although superficially similar to the genitive sentences discussed by Lasnik, could not be accounted for by his theory; this formed the basis for the argument that cyclic nodes have no bearing on the issue. The fact that (70) and related structures are all predicted to allow coreference is encouraging for c-command, despite the matter of Langacker's disagreement with the judgment. One potentially serious objection to the constraint is the sheer number of questionable judgments to be found in Reinhart's work. Readers who have only a passing familiarity with c-command are invited to decide which of the following of Reinhart's sentences (some of them adapted) are free, and which are blocked:

71. For Irene's husband, she would give her life.

72. Fond of Donna's car though she is, I like it even more.

73. In Ben's film of Rosa, she found a scratch.

74. For her new glasses, how much did she pay?

75. In John's apartment, he smokes pot.

76. In Carter's home town he spends his most pleasant hours.

77. Michael's problems, he won't talk about.

78. When Lorraine finishes school, she promised Jim to go to Paris.

Some readers will be surprised to learn that all of these are starred by Reinhart. In a brief survey of naive speakers, I found such massive disagreement, not only with Reinhart but among subjects, that

either the theory or the tradition of collecting speakers' judgments must be thrown into question. This matter is discussed in deeper detail in Chapter IV.

2.11 Chomsky's binding conditions

Chomsky (1980, 1981) has adopted the c-command restriction in formulating his three Binding Conditions, which can be paraphrased as follows:

(i) A reflexive or recriprocal must be coreferential with any NP c-commanding it within the minimal S or NP that contains a V, A, N, P, or TENSE node which c-commands the reflexive or reciprocal.

(ii) A personal pronoun may not be coreferential with any NP meeting the conditions outlined in (i).

(iii) Lexical NP's may not be coreferential with any other c-commanding NP.

Conditions (ii) and (iii) are most relevant for our purposes. Consider the effect of these conditions in the following familiar examples:

79. *John hates him.

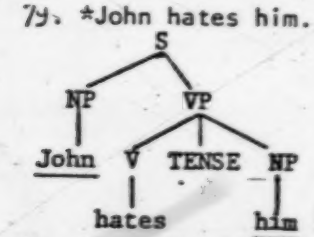

where John c-commands him within S which contains TENSE, which c-commands him.

80. *Near Dan, he saw a snake.

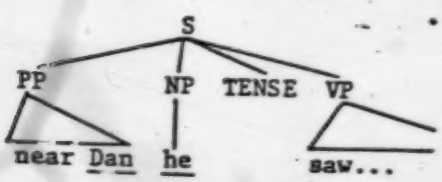

where he c-commands Dan within S, which contains P, which c-commands Dan.

81. *He said that John was sick.

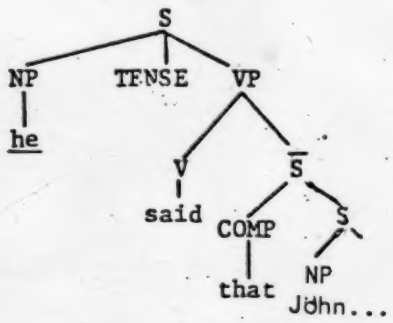

where no NP c-commands he within S; John may not be coreferential with he, which c-commands it.

82. John said that he was sick.

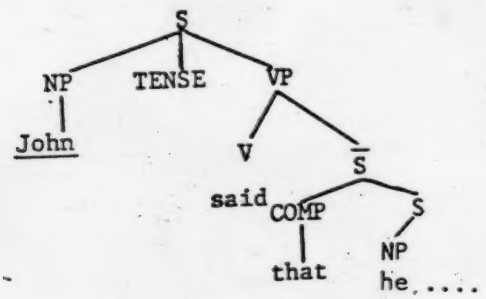

no NP c-commands John; condition (iii) does not apply. The minimal S containing TENSE, wich c-commands he, is S; no NP c-commands he within S.

83. The man he saw downtown gave Harry a gift.

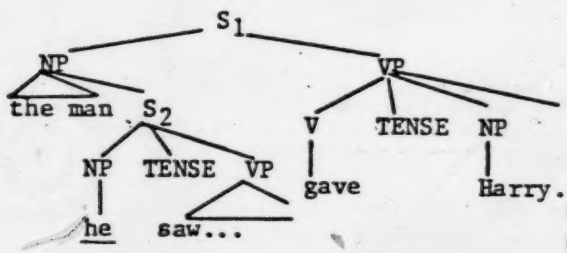

no constraints apply: neither Harry nor he c-commands the other by the conditions in (i).

Thus the binding conditions cover the same range of data as Reinhart's original analysis. One significant difference is that condition (ii) accounts for the Clausemate Constraint as well as other cases of blocked anaphora; in Reinhart's model, the Clausemate Constraint was not accounted for. On the other hand, both conditions (ii) and (iii) are

necessary to account for the full range of blocked and free cases - there is no single interpretive principle.

In subsequent discussion, it is therefore assumed that difficulties in interpreting the c-command constraint in performance terms apply equally to the interpretation of binding conditions.

2.12 C-command, adverbial scope, and process interpretations

One of the most notorious problems with c-command involves its handling of preposed adverbials. In this section I shall attempt to shed some light on this issue by considering how c-command might be implemented as a coreference decision routine in a machine natural language parser. The purpose of this exercise is not to build a straw-man argument against c-command by demonstrating its failure to model human comprehension processes. On the contrary, in addition to helping clarify certain questions of syntax versus semantics, the machine implementation problem will provide a useful example of the very different goals of product and process models.

Reinhart notes sentence pairs purported to differ in coreference options because of differences in the VP which give rise to differences in adverbial scope. Consider these examples:

84. In Ben's office he is an absolute dictator.

85. * In Ben's office he placed his new brass bed.

Although Wasow (1972) and Bickerton (1975) have suggested that the differences are attributable to performance constraints or presupposition respectively, Reinhart claims that the differences are due to purely structural factors: the PP of (84) is attached at S, while

that of (85) is attached at VP. In surface structure, this is represented by having the VP-scope adverbial fill the COMP position, while the sentence scope adverbial is attached to a higher S node:

84. In Ben's office he is a real dictator.

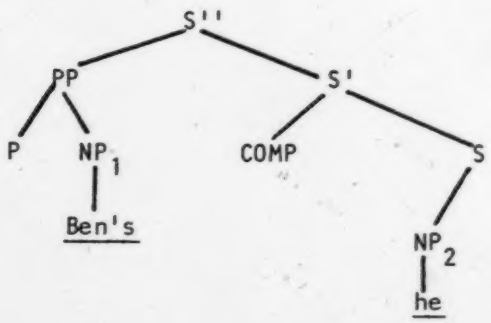

85. * In Ben's office he placed his new brass bed.

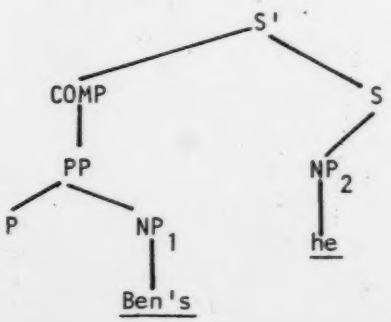

Since c-command may involve only one higher node of the same category, c-command prevents coreference only in (85).

If a language processor is to apply c-command, it must have as input surface phrase markers such as (84) and (85), or construct them from the incoming speech signal. Notice that in order to be able to determine which type of analysis is appropriate, the processor must

have access at least to the verb: for be, S-attachment is correct, while for place, VP-attachment is correct. This then determines the placement of COMP, which is the decisive factor in allowing or blocking coreference, according to Reinhart.

However, this seems an inefficient routine for a parser. A simpler solution would be to determine coreference options directly, by referring to the following interpretive rule:

86. A pronoun may not be coreferential with an NP in a preceding PP of sentence scope.

This obviates the need to work out the c-command domains. Why, then, does Reinhart choose the more complicated route?

The answer lies in Reinhart's contention that the COMP analysis is purely structural, and her desire to provide a unified model of coreference. In order to do so, she must demonstrate that the scope information is not in fact a semantic distinction by providing syntactic evidence that the two PP types are different. The evidence involves the following examples of Though Movement (87) and Pseudo-cleft Formation (88):

87. An absolute dictator though Ben is in his office...
*Place his new brass bed though Ben did in his office...

88. What Ben is in his office is an absolute dictator.
*What Ben did in his office is place his new brass bed.

These pairs show that each rule must operate on an entire VP; when the PP is extracted from the VP, ungrammatical sentences result. These can be remedied by maintaining the original VP structure:

89. Place his new brass bed in his office though Ben did, he...

90. What Ben did is place his new brass bed in his office.

This is a standard form of syntactic argumentation, but I do not

find it compelling; it is simply an attempt to save the analysis from splintering into separate syntactic and semantic components. In fact, one might equally argue that the restrictions on Though Movement and Pseudocleft Formation are merely syntactic reflexes of the semantic scope distinction. In both cases, the claim can be understood in two rather different ways. First, in Reinhart's view, it is COMP position that determines the scope interpretation, in accordance with the notion that structures are first generated by the grammar, then assigned semantic interpretations. Secondly, from the parser's standpoint, COMP position is merely a redundant coding of scope, which can be interpreted directly from (86). Reinhart is free to choose the more computationally costly route in order to have c-command interpret coreference in all cases. A parser, on the other hand, can omit the c-command analysis.

2.13 Resolution biases and processing strategies

It often comes as a surprise to the non-linguist that syntactic rules governing anaphora do not state which NP goes with which pronoun. That is, the syntactic constraints are not deterministic resolution principles: while some variant of command theory may specify which NP's may not be coreferential with a given pronoun, one cannot determine the antecedent actually intended by the speaker on the basis of syntax alone.

Given the interest in syntactic processing strategies sparked by Bever (1970) and Kimball (1973), it is therefore somewhat surprising that the issue of preferred antecedence has not been seriously con-

sidered in the linguistic or psycholinguistic literature. That such biases do exist can be surmised from the fact that sentences such as (91) and (92) are not normally perceived as ambiguous:

91. In the town where Doug decided to live with Rory, he is required to pay a special gun tax.

92. When John gave the money to Jeff, he glanced at George.

in spite of the fact that in each case the alternative NP may function as the antecedent:

93. In the town where Julie decided to live with Rory, he is required to pay a special gun tax.

94. When Sue gave the money to Jeff, he glanced at George.

The syntactic, discourse, and pragmatic basis of such biases has not, to my knowledge, been carried further than Jackendoff's (1972) mention of the subject resolution bias for reflexives such as:

95. The waitress showed Sheila a picture of herself.

(see page 12 above). In this section I shall sketch the kinds of variables that seem relevant to such research.

First, the relation of an NP to the verb appears to affect access to an ambiguous pronoun: subjects, as Jackendoff noted, are more available than non-subjects, even when the pronoun is in a higher sentence, as in (93) and (94). However, this is not always the case:

96. Fred gave a book by John's father to his friend.

97. Ellen was seen by Mary near her house.

98. Ellen saw Mary near her house.

99. Rosa won't like Mary anymore, with her mother interfering all the time.

100. Bill broke Dan's new glasses in his apartment.

101. Bob showed Bill the microphone that was hidden in his office.

In these cases, a genitive pronoun in a PP is equally likely to be coreferential with the preceding subject and with objects, genitive NP's modifying objects, and object agents. It seems too that any slight bias is attributable to pragmatic, rather than syntactic, variables: Resolution to Bob seems better for (101) if one considers it more likely that a person would reveal a bug in his own office to someone else, but resolution to Bill seems better if one thinks it more likely that a person would reveal a bugging to its victim.

In still other cases, the bias involves the extent to which a sentence is more naturally interpreted with internal resolution. The following two cases suggest that external resolution is sometimes preferable:

102. We'll just have to fire him, whether Smith knows about Jones or not.

103. We parked her car near Mary's house.

Conversely, sentences for which internal resolution is quite normal:

104. John said that he was sick.

are easily given external resolution by providing a topic NP:

105. George arrived today. John said that he was sick.

Judgments for a large variety of cases of free anaphora would provide an initial view of the strategies speakers use in interpreting these potential ambiguities; analysis of spontaneous speech samples would in all likelihood prove even more fruitful.

2.14 Anaphora and discourse structure

2.14.1 Anaphora without syntax

Despite the fact that anaphora is widely recognized to be a cen-

tral element of discourse cohesion (e.g., Halliday and Hasan, 1976), relatively little has been written about the effects of discourse structure on either anaphora resolution or constraints on intersentential coreference. There is, however, a school of thought which proposes to account for constraints on anaphora, as well as the relaxing of certain constraints, partly or entirely in terms of discourse structure.

Among the most radical proponents of this view are Keenan (1976) and Bolinger (1977,1979), whose general attitude is that there are no independent syntactic effects. Bolinger cites approvingly Keenan's comment that there are "no structurally statable restrictions", and Bily's (1977) claim that attempts at improving precede-and-command theory "are often mere ad hoc patchworks on the basic syntactic rule which does not work" (cited in Bolinger, 1979).

On closer examination, however, Bolinger's work is seen to have little relevance to the general issue of constraints on coreference; rather, it is directed toward the very different question of when in discourse speakers choose to use a pronoun over a lexical NP. Only occasionally does this affect the grammaticality of a given coreference reading.

Consider the case (first noted by Kuroda, 1968) of disallowed backward pronominalization for indefinite NP's as opposed to permissible coreference with definite NP's:

106. When he entered the room Mary saw (*a salesman) (John) smile.

In addition to pointing out that this is not statable as a syntactic constraint, Bolinger notes a context in which the condition is relaxed:

107(a) I guess I'm enough of a pushover to take any woman who

was reasonably willing. -- If she asked you, would you marry a widow with five children?

However, I believe that it is the conditional, and not the context, that relaxes the constraint:

107(b) *When/If she came to my door I told/would tell a prostitute to leave.

Note that substituting beggar for prostitute reveals further pragmatic effects. Furthermore, the contrast putatively exhibited in (106) was not found in the judgment data of Chapter 3 below.

Stronger claims about the role of discourse are advanced by Delisle (1973) (and later adopted by Kuno, 1975), who claims that free backward pronominalization, which later motivated the development of c-command theory, does not involve a rule of sentence grammar. His position is essentially that in sentences such as

108. The woman who is to marry him will visit Ralph tomorrow.

the pronoun refers not to the following Ralph, but to an antecedent in prior discourse. If this is true, it is unnecessary to account for backward anaphora at all. Instead, sentences such as (108) would be instances of forward discourse anaphora, and the full NP Ralph must be viewed as nothing more than a resurrection of the already-resolved antecedent.

For such sentences, Delisle claims that an antecedent must be established in prior discourse. There are several arguments against this view. First, Carden (1978) cites many counter-examples:

109. While he hadn't read the article, Associate Dean...Martin Gifford said...

Fully 51% of the instances of backward pronominalization in texts surveyed by Carden involved reference which could not be predicted

from prior context. Secondly, neutral contexts can be found in which intra-sentential resolution is obviously free and backward:

110.(a) Did you hear the news? (b) Did you hear Fred?
(c) The woman who is going to marry him will visit Ralph tomorrow.

And finally, the argument that discourse-forward anaphora excludes the possibility of backward is not only non-demonstrative, since both forward and backward anaphora might be involved, but also somewhat ludicrous when one then considers what coreference relation obtains between the pronoun and the later full NP: surely coreference is not blocked (as shown by 110a-c), but if it is not free, then it must be obligatory-forward, like reflexives. This cannot account for the ambiguity of 110b-c.

2.14.2 Kuno's discourse constraints

Kuno (1972 a,b; 1975) attempts to demonstrate that pronominal anaphora and reflexivization are "controlled primarily by nonsyntactic factors" (1975, p. 276). Here we shall restrict discussion to his comments on simple pronouns.

Unlike the majority of his contemporaries, Kuno's approach is transformational rather than interpretive; he assumes that pronominalization is a rule "which applies after all movement transformations apply" (p. 280), but views it as a replacement rule. Rather than specifying the conditions under which a pronoun and a lexical item may be coreferential, Kuno prefers to view the matter in terms of the discourse conditions under which a particular NP may be replaced by a pronoun. In this respect, his goals are similar to those of Bolinger.

Kuno's principles are divided between Discourse and Sentential factors. The most general of the discourse rules is the requirement that all NP's with referents which are "clear from the preceding discourse" (p. 280) are to be pronominalized.

Although Delisle (1973) made a similar assumption, it is not a necessary one. For example, reidentification is a frequent device, both in discourse and in single sentences:

111. Why would you give Oksana a pair of pink Pierre Cardin socks?
Because Oksana likes pink socks!

112. Whatever Lola wants, Lola gets.

In general, then, this constraint cannot explain coreference between lexical NP's.

Kuno's second principle is that if an NP has a discourse antecedent, all but one occurrence of the NP should be pronominalized, subject to a variety of conditions, including the precede-and-command constraint and the predictability requirement. Although this is at odds with Carden's demonstration that unpredictable backward pronominalization is also possible, Kuno's constraint might be viewed as useful for just those cases in which a sentence with free backward anaphora has no antecedent in prior discourse. The constraint prohibits pronominalization of the leftmost NP if its antecedent is not predictable:

113. Who did you calm when you saw him getting mad?
#I calmed him before Henry did something rash.
I calmed Henry before he did something rash.
(p. 283)

This implies that discourses with explicit mention of the antecedent are better; however, they seem to suffer from a lack of motivation

for reidentifying the second NP:

114. What did you do when you saw Henry getting mad?
?I calmed him before Henry did something rash.
I calmed him before he did something rash.

Moreover, the generality of the constraint is uncertain. I strongly disagree with Kuno's judgments in cases such as:

115. What should I do next?
*Go to your house, and if you can lift it, bring back the big rock...

116. What should I do next?
*If you can find him, ask John.

Not only are these acceptable pronominalizations, but they sometimes break Kuroda's indefiniteness constraint, despite Kuno's claim that accounting for indefiniteness effects is one of the advantages of his own approach:

117. Before I could stop it, a huge boulder rolled over the edge and crushed my new bike.

Furthermore, some of Kuno's own examples serve to demonstrate that for certain apparent violations of his constraint, sentence-internal coreference is indeed quite possible:

118. What do people eat in their lifetime?
By the time he reaches the age of 70, the average American consumes...

Perhaps Kuno's objection to this discourse stems from the fact that it uses a pronoun to shift the topic from people to the average American. In other words, there appears to exist a weak discourse constraint against using backward anaphora to introduce a new topic, since the topic pronoun will normally be taken to refer to the current topic. This is evident in (119), in which a topic shift, even if intended by the speaker, cannot be inferred by the listener:

119. Fred sold me a car last week. When he comes over this afternoon, tell John to give me a call.

Such a constraint is well motivated, but has little to do with Kuno's point; in his example, forward pronominalization is neither better nor worse:

120. What do people eat in their lifetime?
By the time the average American reaches the age of 70, he consumes...

Yet the fact that topic shift lies at the heart of the problem simply serves to demonstrate that people and the average American are not clearly related; in order to see the anaphoric relation, one must suppress the topic shift and infer that people and the average American are intended to be coreferential. To clarify this relation, one has only to substitute Americans for people in the original example:

121. What do Americans eat in a lifetime?
By the time he reaches the age of 70, the average American consumes...

Thus it would seem that Kuno's second principle is quite useless.

Another of Kuno's discourse constraints is the Constraint on Predictable Theme Pronominalization, which states that Pronominalization cannot apply when the to-be-pronominalized NP is both the theme of its sentence and predictable from prior context:

122. #If John can, he will go to see a movie.

123. If he can, John will go to see a movie.

In (123), for example, the subject of the matrix clause is thematic and predictable in that the prior question refers to John. However, Kuno's constraint seems to apply to the wrong NP: what is required is not a prohibition against a thematic pronoun, but one against a non-thematic full NP:

124. What will John do this Sunday?
If he can, he will go to see a movie.

In (124), the discourse breaks the predictable theme principle, but is much better than (122). While increasing the distance between the initial NP and the later pronoun improves (122):

125. If he can, that is, if Marsha doesn't make him paint the garage, John will see a movie.

no manipulation will make the full nominal natural in the subordinate clause of (122).

At any rate, these facts are better accounted for by a separate condition, also proposed by Kuno, which disallows pronominalization in the "known" portion of a backward pronominalization sentence:

126. Did you see John when he came to Boston?

(a) *Yes, I saw him when John came to Boston.
(b) Yes, I saw John when he came to Boston.
(c) Yes, I saw him when he came to Boston.

Notice, however, that it is more natural to use two pronouns, as in (126c).

It bears mentioning that (126a) is one of the disputed structures starred by Langacker but accepted by Reinhart (see page 22 above). Given the inappropriateness of this structure in a context where John is already "known", and the general principle that backward pronominalization will be interpreted as exophoric in discourse, it is not surprising that such sentences will often be judged ungrammatical. The psycholinguistic processes underlying such judgments are discussed in Chapters III and IV below.

To summarize, Kuno's arguments concerning the discourse basis of certain anaphoric processes are unconvincing, often based on questionable

judgments, and, although raising interesting questions about the role of discourse principles in anaphora, do not succeed in showing that anaphora is "controlled primarily by nonsyntactic factors" (p. 276).

2.15 Discourse cohesion effects: Relevance and topic shifts

In this section I consider the effect of discourse cohesion on anaphora interpretation. Consider these discourses:

127. Ann got home late last night.
Ethel says she was in an accident.

128. The doctor just arrived.
Gary is afraid he is having one of his attacks.

How can one explain the fact that anaphora is resolved to the topic of (127), but sentence-internally in (128)?

The answer seems to be that interpreting she as Ethel in (127) would result in an incoherent discourse, in which the speaker first mentions Ann's arrival time, then irrelevantly adds that Ethel had an accident. In (128), conversely, one infers that it is Gary who is having an attack, thus explaining the presence of the doctor. The alternative resolution, in which a doctor arrives and has an attack, is dismissed as unlikely.

Notice, however, that while resolution appears to be determined by a topic resolution strategy in (127), the more natural resolution in (128) involves a topic shift. Clearly, then, discourse topicalization is not the only extra-sentential factor governing resolution.

These examples demonstrate that anaphora itself is a vital component of topic maintenance and topic shift. Consider the following discourses, in which an error in interpreting the speaker's intended resolution also results in an error in topic interpretation:

129. Ann got home late last night. Ethel says she was in an accident, so that when she opened the door, she didn't notice that it was Ann standing there.

130. The doctor just arrived. Gary is afraid he is having one of his attacks, and that's why he forgot to call before coming over.

Even when one attempts to construct an incoherent discourse with potential inter-sentential anaphora, the reader attempts to find a means of determining either a reason for the topic shift, or an interpretation under which the topic is maintained:

131. Ann loves chocolate cake.
Ethel says she was in an accident.

It is difficult to read this discourse without speculating that Ann's love of chocolate cake somehow resulted from an accident, or that her love of chocolate cake led to the accident, or that Ann's love of chocolate cake resulted in Ethel's accident. The least likely interpretation is that the speaker has spontaneously uttered two completely unrelated sentences.

The mechanism underlying these observations is Grice's (1975) conversational maxim, "Be relevant". The search for relevance is an essential ground rule in co-operative conversation, and is crucial to the description of how shared knowledge affects the structure of linguistic interactions (Smith, 1982; Sperber and Wilson, 1982). In (128) and (130), the listener also makes use of world knowledge about doctors and house calls in order to establish coreference and topic structure. The modelling of such <u>scripts</u> is an important goal of both cognitive psychology and artificial intelligence (e.g. Schank and Abelson, 1977).

In summary, then, topic maintenance and topic shift are not simply factors which help to resolve anaphora; rather, they are often determined by how anaphora is resolved in a given context.

2.16 Scene shifts

Chafe (1976) has hypothesized that items leave a listener's consciousness when the conversational focus shifts from one scene to another, and that a co-operative speaker no longer refers to such items anaphorically (or with reduced stress associated with givenness) after a scene shift is realized. Webber (1979) provides this example:

> 132. Fred left his niece at home and headed for the zoo with Mary and John...When they got to the zoo,...the girl... (p. 1-7)

Here, the fact that the niece does not participate in the zoo scene means, in Chafe's terms, that Mary is the only feminine NP currently in the listener's consciousness, such that definite reference may be used unambiguously (the girl); in fact, it even allows the reader to infer that Mary is relatively young.

However, it is unlikely that this phenomenon has anything to do with the psychological notion of consciousness, which is normally taken to mean, roughly, that a concept is currently active in short term (working) memory. Chafe circularly defines consciousness as the domain of entities which can be referred to in this way, but it can readily be demonstrated that scene shifts do not necessarily eliminate NP's from the former scene as antecedents for a pronoun mentioned after a scene shift has taken place:

133. ... and that's why Bob had decided to stay behind in New York. Meanwhile, off in Cairo, Greg was writing to him every day.

If Bob were indeed no longer retrievable, the full NP would be necessary in the second sentence. Moreover, it is possible for the scene shift itself to invite the inference that the antecedent must be in a prior scene:

134. ...and that's why Bob had decided to stay behind while Fred and Greg taught Arabic at the American school in Cairo. So a month later in Cairo, Fred was writing to him every day.

In this case the pragmatics of the scene shift is unrelated to which scene is currently foregrounded, and the effect has nothing to do with consciousness, except in the circular sense adopted by Chafe.

The notion of scene shift as a discourse factor contributing to anaphora resolution is perhaps best viewed as only one of a very large number of inference types that can affect the process. Chafe's attempt to explain it in terms of memory processes, however, is unsuccessful.

2.17 Discourse referents and discourse model synthesis

Karttunen (1969, 1976) examined the conditions under which an indefinite NP establishes a "discourse referent", that is, an antecedent whose existence may be inferred from the very fact that an anaphoric expression is used. He pointed out that this notion explains why an NP cannot always be referred to anaphorically:

135. Bill doesn't have a car. *It is black.

Such constraints on pronominalization are reminiscent of Postal's (1969) "anaphoric islands":

136. Max is an orphan, but MY parents are alive.

137. *Max is an orphan, but MINE are alive.

where the implied NP the parents of Max can be contrasted with my parents, but not with the elliptical mine. Hankamer and Sag (1976) show that similar restrictions apply when referents are situationally rather than linguistically available:

138. (Hankamer attempts to stuff a 9-inch ball through a 6-inch hoop)
Sag: #It's not clear that you'll be able to.
It's not clear that you'll be able to do it.
(p. 392)

Webber (1979) approaches this issue in terms of what the discourse "makes available" (p. 1-15) as antecedents for pronouns: antecedents are very often not explicit NP's, but entities inferred from the ongoing construction of a discourse model. The following are only a few of the examples she provides:

139. Blend a cup of flour with some butter. Moisten it with some milk (it = the flour-butter mixture), then knead it (it = the flour, butter and milk mixture) into a ball.

140. A Rhodesian ridgeback bit me yesterday. They are really vicious. (they = the class of Rhodesian ridgebacks.)

141. John dunked Mary's braids in the inkwell. It made her cry. (it = the specific event of John dunking Mary's braids in the inkwell).

142. John dunked Mary's braids in the inkwell. It always makes her cry. (It = any occasion on which John dunks Mary's braids in the inkwell).

Notice in particular that the pronoun does not necessarily even agree morphologically with any overt NP (e.g., 140), and that the antecedent need not even be an NP (e.g., 141, 142).

Finally, Karttunen demonstrates that an antecedent may exist not in the speaker's real world, but in a possible world created via the presuppositional structure of certain verbs:

143. John regrets that he doesn't have a nicer car. He would keep it in a heated garage. / *He keeps it in a heated garage.

The notion of discourse model synthesis (Webber, 1981) as an explanation for the generation of non-overt antecedents is an important technical advance for both psycholinguistic and machine comprehension models.

2.18 Summary of formal treatments of anaphora

In the foregoing review of formal analyses of pronominal anaphora, I have emphasized the diverse nature of the phenomenon. Despite attempts to explain anaphora entirely in terms of resolution biases, discourse structure, pragmatic factors, and other higher-level effects, it seems fair to say that the original question, the principle or principles governing the distinction between free and blocked anaphora at the sentence level, remains unanswered. None of the sentence-level theories is completely satisfactory, and although there is general agreement that the distinction is related to depth of embedding, it is far from clear that all such distinctions can be subsumed under a single interpretive rule, or even that such a rule must make reference to syntactic structure alone.

In Chapter III we shall consider some of the processing factors which affect coreference judgments. As pointed out in 2.10, coreference judgments are far from clear, and this is probably the most important stumbling block to be overcome in answering the fundamental questions about the syntax and semantics of anaphora.

CHAPTER THREE

COREFERENCE JUDGMENTS, COMPREHENSION STRATEGIES, AND BLOCKED ANAPHORA

3.1 Introduction

In Chapter II we examined various accounts of the syntactic basis for native speakers' judgments of the acceptability of coreference, and the problem of inconsistent responses to the judgment task, both among linguists and among naive speakers, was mentioned in passing. In this chapter we shall consider some of the complications which arise when one attempts to account in detail for these unexpected deviations from the predictions of formal models.

These deviations can be attributed to several factors, not the least of which is the fact that in the anaphora literature, as in many other areas of linguistic investigation, linguists' judgments with respect to the crucial sentences are not representative of the intuitions of naive speakers. In addition, we shall see that the task of judging coreference possibilities is not a direct window on speakers' intuitions; rather, judgment presupposes comprehension, and in the case of the unusual sentences which are used to test the formal theories, a wide variety of processing strategies may be used - strategies which

do not necessarily reflect the speaker's internalized grammar. Evidence for the existence of such strategies will be presented in this chapter, while Chapter IV reports a judgment experiment which confirms the hypothesis that coreference judgments are heavily influenced by analogy and contextualization effects. Ultimately, it will be argued that judgment data, both from linguists and from non-linguists, cannot be accepted as a reliable indication of the mental representation of grammatical knowledge in this domain, and that only production data (such as text counts and analysis of recorded speech samples) can give proper insight into the nature of coreference constraints.

The structure of this chapter is as follows. In 3.2 we shall consider some of the variables which can affect a listener's decision as to whether or not coreference is possible in a particular sentence. In the course of this initial discussion, I report the results of a survey of a small number of native speakers which suggests that a surprising number of blocked cases are normally understood as allowing coreference. Careful analysis reveals plausible processing explanations for this result. In 3.3 we focus on methodological problems inherent in the gathering of coreference judgment data, with particular reference to a paper by Carden and Dieterich (1980). This raises the question of how "contextualization", both explicit and implicit, can influence judgments, and leads to the adoption of the experimental technique used in the study reported in Chapter IV. The last section, 3.5, attempts to place the results of this study into perspective with respect to the use of judgment and comprehension data in linguistic investigations.

3.2 Variables affecting coreference judgments

3.2.1 Assumed stress

Akmajian and Jackendoff (1970) provide several examples of how syntactic coreference options may be altered by changes in stress:

144. After he woke up, John went home.

145. *After he woke up, JOHN went home.

146. After HE woke up, John went home.

I do not find these examples particularly convincing: apart from the intended noncoreferential readings, the same stress pattern may also indicate contradiction under coreference:

147. Did you say that BILL went home after he woke up?
No, I said that after he woke up, JOHN went home.

148. Did you say that Lee went home after SHE woke up?
No, I said that after HE woke up, Lee went home.

Since this stress pattern allows both the coreferential and the non-coreferential readings, it must still be considered a case of free anaphora.

Nonetheless, there do exist cases in which a change in NP stress brings about a real change in coreference options. Consider (149):

149. They parked his car near Ian's house.

This is normally interpreted as blocked. However, some speakers are unwilling to state that coreference is impossible, and I believe that this is related to the fact that while stressing Ian results in definite blocking, stress on house makes coreference possible:

150. *They parked his car near IAN's house.

151. They parked his car near Ian's HOUSE.

It is important to note that if sentences are presented in written

form (e.g., 149), subjects are free to interpret the stress pattern in whatever way they see fit. This could result in a judgment of absolute non-coreference, despite the fact that the subject might accept the same sentence as anaphoric had he or she thought of the alternative stress pattern.

This seriously undermines the notion that the essential predictor of anaphoric relations is syntactic form, and argues for the more complex view that syntactic and suprasegmental features operate as independent but interacting parameters in coreference assignment. The example under consideration is particularly enlightening in this regard, since even appeal to "normal stress" (see Schmerling, 1974, for a critique of this notion) cannot resolve the issue. The "normal" stress for (149) is (150), with stress on the possessive lexical NP, and no possibility of coreference. The precede-and-command analysis predicts that coreference should be blocked in both cases, while c-command predicts that coreference should be possible for both. Neither theory can capture the fact that the coreference interpretation depends on the interaction between this syntactic form and stress. Since each theory is wrong for one of the stress versions, it is impossible to decide whether stress should be viewed as relaxing the blocking rule for (151), or as blocking otherwise free coreference in (150).

3.2.2 Length

Akmajian and Jackendoff also claim that as a sentence is lengthened, a coreferential reading often becomes easier to obtain:

152. *In John's apartment, he smokes pot.

153. In John's brand new apartment on Scrae Street, he smokes pot.

They attribute this to the fact that

> as the sentence becomes longer, the relative stress levels of the NP's and pronouns decrease with respect to the clause containing them. (p. 126)

Lakoff (1968), Reinhart (1981) and Wasow (1972) note similar effects.

While several speakers whom I surveyed agreed that (153) is somehow more grammatical than (152), none objected to coreference in the latter; two of the judges corrected the sentence to "John smokes pot in his apartment", but this stylistic alteration should be viewed not as rejection of the coreferential reading, but only as a judgment that this means of expressing coreference is marked.

Furthermore, there is no evidence at all to support the notion that the NP's concerned are somehow destressed. Even if careful acoustic analysis were to find evidence of destressing, it is not clear how an imperceptible stress difference could have the same effect on semantic interpretation as the contrastive stress differences in (144) - (151).

Reinhart (1981) attributes the effect not to stress, but to one's perception of the "relative importance" of the PP as the sentence is lengthened: the longer PP is less a locative verbal complement than an independent proposition, and the matrix clause becomes a statement about the PP.

Of course, Reinhart's suggestion is ad hoc. She does not provide further evidence to support the notion that length and "importance" are correlated, nor does her hypothesis make quantitative

predictions about this correlation: how many words or constituents are needed to bring about a change in coreference? For example, O'Grady (1982) shows that the addition of only one word can bring about a significant change in interpretation:

154. ?His mother likes John.

155. Only his mother likes John.

In addition to the lack of agreement among speakers as to whether (152) is blocked, none of these explanations accounts for the improvement in (157) over (156):

156. *John told her that Meryl was lazy.

157. One of the members of the Academy seems to have told her that Meryl may be up for an Oscar.

Reinhart's analysis fails here: there is no possible scope explanation. Akmajian and Jakendoff's explanation also fails, since the same sequence (told her that Meryl) may be pronounced with identical stress, rate and intonation, yet the difference in the acceptibility of coreference remains.

It would appear, then, that all we can say with certainty about length and coreference is that a longer version of a sentence sometimes allows coreference, while a shorter version does not. This is somewhat reminiscent of Ross's (1967) Heavy NP Constraint, in that it invites a processing solution (e.g. an unspecified memory constraint) which cannot be fully explained by appeal to a specific processing theory.

Such evidence poses a difficult problem for both grammatical and psycholinguistic accounts of anaphora: should either type of model block coreference for the particular dominance relation involved, and

relax the constraint under poorly defined length conditions? Or should a theory allow coreference (thereby vitiating current syntactic explanations) and block it only when the sentence is "too short"? All of the above-mentioned theorists have adopted the former solution, but the fact that at least some speakers accept coreference even with the shorter versions suggests that the latter approach may be a better one. It is also possible that speakers who accept coreference in the short versions do so on analogy to the longer versions, and that although they would always interpret such a structure as blocked during normal processing, the self-conscious metalinguistic task of coreference judgment encourages the use of such analogical processes.

It should be obvious, however, that appeals to native speaker intuitions cannot be operationalized in a judgment task. Perhaps the only way to avoid the circularity of this argument is to abandon the notion that speakers' judgments can reveal their grammatical knowledge, and accept only production data as reliable. One could accept the original formulation of the length effect if preposed PP's such as that of (152) are very rare or non-existent. Lack of production of such forms is much better evidence of the operation of a constraint than inconsistent responses from confused subjects to sentence types they would not use.

3.2.3 Perceived dialect variation

Subjects performing a judgment task may be influenced by their impression that a particular form is characteristic of a different

social or regional dialect. One of the subjects I surveyed based her coreference judgment for (152) on her impression that Yiddish-influenced dialects might allow such structures. However, she was uncertain as to whether coreference would indeed be possible in such cases, since she was not in fact very familiar with the dialect. Householder (1973) also mentions the possibility that grammaticality judgments can be influenced by the speaker's perceptions of other dialects.

3.2.4 Connectives

Bolinger (1979) claims that syntactic accounts are weakened by the fact that the type of connective used between clauses can alter coreference possibilities:

158. She could pass for my sister, though Joan isn't related to me at all.

159. *He had promised to help me, but John refused.

This argument is simply wrong. First, by the simple expedient of substituting but for though in (158) and (159), we see that coreference is unaffected by the connectives:

160. She could pass for my sister, but Joan isn't related to me at all.

161. *He had promised to help me, though John refused.

Secondly, the fact that coreference is blocked in (159) is hardly an argument that anaphora is blocked in the sense intended by other linguists. Coreference is not possible in this particular token, but the syntactic form does allow coreference in other cases. Lack of coreference between John and he can be attributed to pragmatic factors; under the coreferential reading, (159) is contradictory, and one thus assumes that coreference is not intended. Note that coreference can

readily be obtained if the contradiction is removed:

162. He had promised to help me, but John later refused.

Similarly, coreference can be blocked in (158) for reasons of a semantic and pragmatic nature:

163. *I think she could pass for my sister, though my late friend Joan said that my mother looks too old for that.

These cases are nothing more than specific instances of pragmatic biases in resolving free anaphora. In such cases, Bolinger's use of the asterisk deviates from the practice of other researchers, who define blocked anaphora solely on the basis of syntactic relations.

3.2.5 Pragmatic content

We have just seen that a sentence with syntactically free anaphora may be judged categorically nonanaphoric for reasons unrelated to syntax. This poses another methodological problem for those who claim that native speakers' judgments are an important source of data, in that naive speakers, unable to distinguish between syntactic and pragmatic effects, might "erroneously" judge such sentences as nonanaphoric, firmly denying the possibility of coreference. Variation in the transparency of the syntactically acceptable but pragmatically odd readings of the following examples suggests that interactions between syntax and pragmatics will be difficult to untangle:

164. I'll feed him his favourite dishes every night, whether Rod likes it or not.

165. Robert says that he died in his sleep last night.

166. When Ben went off to Africa to work on that telecommunications project, he stayed behind in Paris.

Although each of the above exhibits free anaphora by all syntactic accounts, the degree of garden pathing seems to be related to the strength of the pragmatic bias, which varies from sentence to sentence (see Hirst and Brill, 1980, for experimental evidence for this type of variation).

There is an interesting interaction between the length factor which relaxes the constraint on backward anaphora with complement clauses:

167. *I told her that Meryl was up for an Oscar.

168. One of the members of the Academy seems to have told her that Meryl is up for an Oscar.

and the pragmatic variation just noted. Although (168) allows coreference, one finds no garden pathing in the following pragmatically biased version:

169. *One of the members of the Academy seems to have told her that Meryl has left for Boston.

This suggests that despite the length factor, the reader is far from committed to the coreferential reading when the last four words are being processed. One might then speculate on why garden pathing does not occur: the coreferential reading for the longer version is marked, these structures are initially processed as non-anaphoric, and a final decision as to the correct reading is put off until all the contextual information has been processed. Such complex linguistic interactions promise to provide a wealth of hypotheses concerning the processing strategies speakers employ during comprehension.

3.2.6 Ungrammaticality due to factors other than coreference

One factor which clouds the linguistic arguments is that sentences which are purported to be nonanaphoric are sometimes starred in the literature for reasons having nothing to do with anaphora, although the asterisk is then interpreted as signalling blocked coreference. Consider Reinhart's judgments on the following, which she considers to be problematic for c-command:

170. *To him, I spoke in Ben's office.

171. *For him, Ben's neighbours gave a party.

Reinhart does not mark stress, but it is reasonable to assume that she intends these to have stress on the preposition, since they are presented in contrast to the following grammatical examples with stress on the prepositions:

172. Behind him, Ben keeps a gun.

173. Under him, Ben's mother found a mango.

I presented the starred sentences to five speakers, who were asked to decide whether coreference could obtain. All five commented on the strangeness of these sentences, presumably because of the fronted dative and benefactive NP's. When forced to make a decision, four of the judges said that coreference would be permissible despite this other strangeness. The fifth refused to speculate, arguing that if she had to imagine some new form of English in which such sentences were possible, she would not be able to predict how coreference might work under the new rules! In attempting to force coreference judgments on otherwise ungrammatical sentences, Reinhart has either been blinded by the fact that no interpretation is normal, or has greatly misjudged

the average person's ability to make semantic sense of asyntactic prose.

This brings to light the problem raised in 3.2.2 above: What do we make of linguists' and naive speakers' coreference judgments on bizarre sentences? Do such sentences really provide crucial information about the generality of an interpretive principle such as c-command? And who is to be trusted - the linguist, so strongly influenced by the perceptual set of his or her own predictions (see Spencer, 1973), or the naive speaker, who may opt out of syntactic processing entirely in order to make a necessary interpretation?

In this case, Reinhart seems willing to assume that a speaker can ignore the problem with PP position, and, like some linguistic automaton, apply the rules for normal cases to the new situation. While the application of a theory to new cases is a venerated scientific procedure, its application to non-cases is a departure from normal practice.

A better interpretation of Reinhart's asterisks is that PP fronting itself is disallowed under coreference for these cases. While this seems to be precisely the problem, it hinges on speakers' acceptance of coreference. Ironically, this is precisely what Reinhart's own theory predicts; had she noticed the true reason for the unacceptability of her starred sentences, she would not have considered the sentences to be problematic.

There is yet another enlightening problem with these examples. If one alters the stress pattern as follows:

174. To HIM, I spoke in Ben's office.

175. For HIM, Ben's neighbours gave a party.

the sentences appear - paradoxically - to be ambiguous between a free

and a blocked reading. That is, individual speakers find it difficult to decide whether the pronoun and the full NP may be coreferential. However, this is an insidious form of uncertainty, since the decision as to whether an anaphoric relation may exist should not be a question of degree (unlike other questions of grammaticality/acceptability), but a binary choice. Logically, either there does or there does not exist a reading under which coreference obtains. In other words, if a speaker decides that coreference is possible, even under very rare conditions, he should not be confused as to whether it might also be impossible.

The only plausible explanation for this kind of confusion is that the sentence is somehow ambiguous between a free and a blocked reading, and that the speaker sees each in turn, changing his judgment with each change in reading.

This is precisely what happens to these sentences when different discourse contexts are considered. In a context in which HIM is both contrastive and exophoric, no anaphoric reading is possible:

176. I hear you really embarrassed George when you criticized him. Did you talk to Dave in front of the whole office too?
No. *To HIM I spoke in Ben's office.

In a context where the contrast cannot involve the prior NP, either the exophoric or the endophoric reading is possible:

177. Nobody even noticed when Mrs. Haynes moved away, but for HIM, Ben's neighbours threw a big party.

178. Sam says that nobody even noticed when Mrs. Haynes moved away, but for HIM, Ben's neighbours threw a big party.

Dinsmore (1979, 1981) has examined the role of implicit contextualization in grammaticality judgments, and argues that a theory derived

from judgment data will be heavily influenced by the (perhaps unconscious) appeal to both linguistic and non-linguistic contexts. This would explain the observed variation in intuitions about these sentences: some speakers should steadfastly deny the possibility of coreference; others should firmly accept it; still others, vaguely aware of the two readings, should waver between the two judgments.

Lastly, and to further complicate matters, analogy may also work across stress types, such that responses to (170) and (171) are based on the same kind of division of opinion found for (174) and (175). For example, Reinhart, who stars (170) - (171), may belong to the group which denies the possibility of coreference in (174) - (175), making implicit reference to the context of (176). My speakers, accepting (170) - (171), should also accept (174) - (175), if the basis for this judgment is the kind of context represented in (177) - (178), and provided that their appeal to context and analogy does not change over time.

In summary, detailed analysis of this example demonstrates that coreference judgments, both by linguists and by naive speakers, do not provide direct access to the speaker's grammar. Rather, they can be strongly influenced by processing principles related to comprehension and contextualization. Again, the extreme complexity of alternative explanations for grammaticality judgments, and the observed variation in speakers' response to the sentences, suggests that such data collection is an unreliable method for accessing speakers' intuitions concerning anaphora.

3.2.7 Analogy and depth of processing

Three of the four speakers in my initial survey allowed coreference on the following sentence, which was starred by Reinhart:

179. *In Ben's film about Rosa, she found a scratch.

One explanation for this judgment is that these speakers interpreted the sentence with the grammatical reading (not mentioned by Reinhart) under which Rosa finds a scratch in some unspecified object. However, all subjects denied having made this interpretation. A related explanation is that instead of perceiving the ambiguity and choosing the less common (but grammatical) reading, the speakers interpreted the sentence as coreferential on analogy with the much more frequent structures not involving picture nouns, and thus not allowing S-scope interpretation:

180. In Ben's letter to Rosa she found a typo.

Further support for the analogical basis of this judgment comes from the fact that even within the category of picture-noun sentences, ambiguities such as that involved in (179) constitute an extremely small range of pragmatic possibilities, since a picture noun will make the PP ambiguous between the locative and the VP-scope interpretations only if the last NP (scratch) can be interpreted as a physical part of the picture noun. Consider this example:

181. In Ben's film about Rosa, she found a husband.

Reinhart's judgment, although understandable when the sentence is carefully dissected, bears no relation to how a given speaker might resolve the anaphora. Given the extreme rarity of conditions under which a sentence such as (179) might be uttered, the speaker does not - or cannot - perform the required syntactic, semantic, pragmatic,

and contextual analysis to sufficient depth. That the confusion is attributable to processing complexity, and not to syntactic form, can be demonstrated by moving the PP to postverbal position:

182. *She found a scratch in Ben's film about Rosa.

This version of the sentence preserves the relevant dominance relations, but is much more readily judged non-anaphoric, since the adverbial is now more readily perceived as attached at VP.

Ultimately, however, one must question the claim that the c-command analysis reflects speakers' judgments, for it makes incorrect claims about the outcome of the judgment process. While it may indeed be true that speakers do not utter sentences such as (179) because of the constraints on coreference which they have internalized, one cannot claim that they will not interpret such sentences as coreferential. Certainly Reinhart did not gather naive speakers' intuitions on these sentences, although she certainly did consult other linguists. What her judgments reflect, however, is not speakers' judgments, but linguists' judgments of how speakers would judge the sentences if they only knew enough about syntax.

Nonetheless, this procedure may not be as far wrong as it may seem. If we consider that linguists base their opinions on an estimate of the likelihood that speakers would produce such a sentence with the coreferential interpretation, and that linguists are probably more aware of what people in fact say, then their opinions may be worth something after all. However, it is only to the extent that linguists are good estimators of production norms that their theorizing can be taken as a valid representation of linguistic representations.

3.2.8 Judgments on isolated sentences: Task demands

The basic premise underlying the use of grammaticality judgments in grammar construction is that native speakers (including linguists) are able to judge the absolute grammaticality of a sentence in isolation. With respect to blocked anaphora, however, this is a particularly risky venture. Even if a speaker would never utter a sentence such as those discussed in the preceding sections without intending external coreference, and even if s/he has never heard or read such a sentence without an appropriate context, it is not entirely clear that the generalizations governing this lack of production are available in the form of a blocking rule which would result in the rejection of anaphoric readings for such sentences presented in isolation. Given the task of judging whether coreference is possible, the speaker must consider the anaphoric reading and decide whether it is disallowed; this is a task of interpretation, and not of well-formedness per se.

Consider this analogy: given the sentence John red book read, a speaker can easily spot the ill-formedness. However, s/he can also interpret the string, and might well respond affirmatively to the question "Can this mean that John read a book?" When testing the acceptability of a given anaphoric interpretation, we also want to know two things: Is the sentence well-formed, given a particular coreference reading? Each of these aspects of the task is important, but each relates to a different aspect of the constraints on anaphora. By asking speakers to judge whether coreference is possible, we are really asking two separate questions; it may be impossible, in principle, to separate them.

3.2.9 Acceptability, preferred readings, and implicit context: More methodological pitfalls

In this section we examine in detail a paper by Carden and Dieterich (1980) which purports to provide a solution to the problem of contradictory judgments in the anaphora literature. However, I will argue first that their attempt to settle such controversies by experimental means is itself fraught with methodological problems, and secondly, that their interpretation of text count data - a technique which I favour - is flawed.

The paper in question deals with a disagreement over the coreference predictions for a sentence type which is a case of blocked backward anaphora according to precede-and-command (S-command), but which c-command would label as free:

183. We'll just have to fire him, whether McIntosh likes it or not.

Reinhart (1981) simply rejects Langacker's (1969) judgment of non-coreference for these structures, and suggests that it is the lack of discourse motivation for backward pronominalization, rather than a constraint on coreference, which accounts for the strangeness of the coreferential reading. Carden and Dieterich (hereafter, CD) suggest that the disagreement needs better empirical investigation before the matter can be resolved, and claim to provide an appropriate methodology for collecting native speaker judgments bearing on the problem.

CD maintain that linguists' judgments on such sentences are suspect, since they may be influenced by what Spencer (1973) has termed "perceptual set", arising from their familiarity with the theo-

retical issues involved. Specifically, Reinhart and Langacker, and their readers, may have been influenced by examples surrounding the crucial one. Langacker contrasted (184) with (185) to demonstrate noncoreference:

184. I knew Harvey when he was a little boy.

185. *I knew him when Harvey was a little boy.

while Reinhart contrasted (183) with (186), starred by both theories:

186. *He was fired, whether McIntosh liked it or not.

Notice as well that Reinhart's pair is intended to illustrate the different coreference options for subjects and non-subjects, one of the proclaimed advantages of c-command theory. If (185) and (186) are in fact found to allow coreference, then c-command loses one of its main claims to superiority. Other instances of possible difficulties with the subject-nonsubject contrast have already been mentioned (p. 27, (65)).

Several problems bear mentioning before we turn to CD's judgment data. First it is not at all clear that Reinhart's judgment on (186) is any less questionable than Langacker's opinion of (185). The fact that both theories predict noncoreference does not mean that this structure is "uncontroversial" to native speakers, who were not consulted on these judgments. This would have been but a minor quibble against CD, had they not later used (186) as a control item against which to measure relative grammaticality in their study. As it turns out, many subjects accepted coreference on this sentence as well!

Claiming that discourse coherence can outweight syntactic factors in grammaticality judgments, CD decided to present the target sentences

to naive speakers for judgment in a neutral context, and to ask them which NP, McIntosh (sentence-internal) or Smith (sentence-external) was the preferred antecedent. On a second pass with the same stimuli subjects were asked whether or not the other NP could be the antecedent.

The target sentences consisted of four types, modelled on the disputed (189) with the pronoun in main clause object position; the "uncontroversially bad" example, with the pronoun in main clause subject position (188); and two uncontroversially good examples, with the pronoun a subordinate clause subject and McIntosh subject or object of the main clause:

187. (CONTEXT) The directors discussed the situation with Smith all afternoon.
188. ("BAD") They finally decided that he would have to report to the new vice president, whether McIntosh liked it or not.
189. ("DISPUTED") They finally decided that they would have to put him under the new vice president, whether McIntosh liked it or not.
190. ("GOOD") They finally decided that, whether he liked it or not, McIntosh would have to report to the new vice president.
191. ("GOOD") They finally decided that, whether he liked it or not, they would have to put McIntosh under the new vice president.

The results showed that (188) patterns with (189), having fewer internal coreference judgments than either (190) or (191). CD therefore conclude that Reinhart was wrong, and that S-command is the better theory.

This conclusion is unwarranted on several grounds. First, it seems significant that both the crucial sentence and the "bad" one received a fair number of positive judgments (7 coreferential, 14 noncoreferential for 188; 7 coreferential, 19 noncoreferential for 189). Indeed, CD had found several instances of both types in a survey of

spontaneous texts, such as this "bad" one:

192. She was sitting with 8,000 people in Madison Square Garden at a New York Apples match, when Phyllis Rothstein of New Rochelle...

CD see their options as follows:

> We may have a good but (accidentally) rare construction; or we may have a bad construction occurring a few times because of errors.

Even in the face of the very large number of coreference judgments given by their subjects and the occurrence of perfectly acceptable examples in published materials, they maintain that since both theories predict (188) to be ungrammatical, acceptability judgments must be in error for both (188) and (189).

Notice, however, that no explanation for this large number of "errors" is provided, and CD do not consider the possibility that the sentences are rare in their samples because they are stylistically marked and characteristic of certain written registers. Perhaps if they had surveyed tabloids such as the National Enquirer or the Weekly World News, they would have found many more examples of this attempt to create suspense syntactically where none exists in the story itself.

On the other hand, if these sentences are perfectly grammatical, why did a majority of subjects not allow coreference? One reason is that by collecting relative grammaticality judgments in context, CD were asking in effect whether coreference was natural in the particular context they had chosen; this has no necessary bearing on whether or not such sentences may allow coreference in other contexts, or in isolation.

Recall that I have provided several examples of how contextualization can bias anaphoric interpretation in free cases such that coreference is blocked in a particular token; this is not an argument that the sentence is syntactically blocked, and CD's evidence would therefore appear to be irrelevant, to the extent that the context they used did in fact bias interpretation in this way. In the case of their "neutral" context paragraph, it should be obvious that the foregrounding of Smith in the context sentence makes this NP more available as the antecedent for the main clause pronoun, in accordance with the principle that free backward anaphora is resolved externally when an appropriate NP is available in the preceding text.

This is, of course, a serious error in experimental design, but it had a rather interesting result: subjects more often resolved extra-sententially for (188) - (189) than for the other sentences. In these cases the pronoun is in the main clause. While preference for external resolution is irrelevant to the question of whether internal coreference is possible, it does raise the possibility that dominance structure is predictive of markedness with respect to internal resolution.

Solan (1978) reports some possibly relevant data. He presented children with four short discourses foregrounding a discourse topic NP (the cow), and followed these by four target sentences:

TOPIC: The cow...

193. The horse told the sheep that he would run around.

194. The horse hit the sheep after he ran around.

195. The horse hit the sheep with his stick.

196. The horse hit the sheep after his run.

His subjects, aged 5 to 7, chose the cow twice as frequently for (193) and (194) than for (195) and (196), although the most frequent responses by far involved internal resolution to horse or sheep. Solan attributes this to a processing strategy based on the fact that the pronoun is in an independent clause in the former cases, but his is a clausemate of horse and sheep in the latter cases. Unfortunately, his data analysis does not include internal coreference biases, which, as pointed out in 2.13 above, have been given far too little attention in the literature. If there is such a bias, however, Solan may be wrong in his general conclusion that children do not use context in resolving anaphora. It is quite likely that children would tend to resolve to the horse, viewing this as a topic shift within the short discourse. In a footnote, he reports that adult subjects choose the cow as the antecedent about 50% of the time, making no distinction between the two within-sentence conditions. Thus only the acquisition data support his notion that clausemate status affects coreference.

To return to the CD study, one might well ask why there were so many judgments of internal coreference, given the possibility of adopting a topic-resolution strategy. The anomaly can be explained by another design flaw: CD used one other passage (in a weak attempt to give their results greater generality), but in this case the context was truly neutral:

197. We looked all over for a suitable speaker. We talked first to John, then to Bruno, and finally to Doug.

In this case, CD succeed in having all three names equally back-

grounded (all are mentioned; none is the subject of the initial sentence). Their target sentences mention John and use the pronoun he in the same structures as (154). In this case we should expect subjects to respond with fewer external resolutions; unfortunately CD do not report their results separately for the two discourses. Nonetheless, it is likely that the large number of coreferential judgments can be attributed mainly to this condition, which offers the least biased test of the internal coreference question. Since overall, subjects accepted coreference on 40-50% of the "BAD" and "DISPUTED" cases, the figures for the neutral condition alone (i.e., 197) very probably exceed 50%. The fact that there is a significant difference in the number of 'yes' responses for the "GOOD" and the "BAD" or "DISPUTED" cases is hardly an argument against the possibility of coreference in the latter cases.

3.3 Conclusions: Judgments, comprehension processes, and language acquisition

In this chapter I have argued repeatedly that the collection of native speakers' judgments about sentences exhibiting putatively blocked anaphora is fraught with methodological dangers. Both non-syntactic linguistic factors (stress, length, dialect variation, semantic bias, ungrammaticality) and linguistic processing strategies (implicit contextualization, analogy, shallow processing) can render such judgments suspect, in that a judgment of possible coreference will not necessarily reflect the speaker's mental representation of the restrictions on anaphora. In other words, native speakers are not good judges of the boundary between free and blocked anaphora.

I have also argued that judgments on nonoccurring sentences should be discarded, and replaced by direct analysis of production data, since the method of collecting judgments is demonstrably not an unbiased procedure for accessing linguistic intuitions underlying performance. That I should suggest a return to the structuralist methodology of supporting grammatical descriptions with data from attested utterances may come as a shock to those for whom explaining linguistic intuitions is the central goal of linguistic theory. However, it is evident that linguists rarely consult native speakers at all, and that in practice, the intuitions described in the literature are only those of linguists. But does this mean that linguists' intuitions provide a better basis for theory construction in linguistics? Certainly not- the number of disputed cases in the experts' own writings, their failure to differentiate between syntactic and supra-segmental, pragmatic, or discourse-related linguistic factors, and their own susceptibility to processing effects, argues quite strongly against accepting their "expert" opinions as any more reliable than those of the native speaker. If linguists' intuitions are to be worth anything, they must reflect the constraints governing production of attested forms; if they do not, those forms themselves must determine the form of the grammar.

Spencer (1973) showed that non-linguists' grammaticality judgments are internally consistent, but different from those of linguists in 50% of the cases. I have claimed in the preceding two chapters that linguists' judgments for anaphora data are similarly non-representative, and that naive speakers should, given the processing explanations

set forth above, accept far more coreference than predicted by linguistic theory. It might be argued that, having accepted that process explanations suggest that judgment data are unreliable as a reflection of underlying competence, an investigation of speakers' reactions to blocked anaphora is pointless. However, there are at least two good reasons why this should be done. First, it is important to establish whether the prediction of very lax criteria for accepting coreference in the judgment task is indeed confirmed. Secondly, studies of the acquisition of constraints on anaphora rely on the assumption that the endpoint of the acquisition process is the adult grammar, as revealed either by linguistic analysis or by adults judgments on the crucial sentence. Since at an early stage children tend to accept coreference in a number of blocked constructions (e.g., Chomsky, 1969; Solan, 1978) this literature is geared toward testing for children's acquisition of the blocking principle. But given the arguments against using either type of judgment data as evidence for the structure of the adult grammar, it would be dangerous to use these as a point of comparison.

Taylor-Brown (1983), for example, reports that adult control responses to eight types of blocked anaphora ranged from 15% to 100% "blocked" responses, and she reports a mean of only 59% "correct" responses from adults. Noting that O'Grady's (1982) computational approach predicts a smaller number of such responses than c-command, she concludes that the former theory is better. However, O'Grady's theory is not a performance model; to the extent that subjects accept coreference because of processing strategies, their "free" responses

are inflated with respect to their internalized system, which is what O'Grady intends to model.

In addition to this general argument for a better understanding of adult judgment process, I would like to point out a glaring methodological error in certain studies which purport to demonstrate the child's acquisition of blocking principles. In many of these studies (e.g., Lust, 1977; Lust, Loveland and Kornet, 1980; Lust and Clifford, 1983; Solan, 1978; Tavakolian, 1977) the children are asked to act out a sentence exhibiting free anaphora when both a toy named in the sentence and another toy are available antecedents. Consider Tavakolian's examples:

198. To kiss the lion would make the duck happy.

199. For him to kiss the lion would make the duck happy.

200. That he kissed the lion would make the duck happy. (p. 176)

In each case, children tended to act out the kissing by moving an animal other than the duck, a result which was interpreted as demonstrating that they blocked coreference. However, this result has nothing to say about the acquisition of a blocking rule; rather, it shows again that for free anaphora, if the discourse or the situational context proved another possible antecedent, backward anaphora is interpreted as exophoric. The type of "blocking" revealed by such studies is not the type of constraint dealt with elsewhere in the literature.

In the next chapter we examine the claim that speakers accept far more coreference than predicted by popular theories of anaphora. If this is so, as predicted by the processing hypotheses laid out above, we

must acknowledge that linguists' judgments are not representative of those of other speakers, and further investigate the validity of the notion that even naive speakers' judgments provide clear access to the structure of their internalized grammars.

CHAPTER FOUR

JUDGMENT EXPERIMENT

Introduction

The process of deciding whether or not an NP may serve as the antecedent of given pronoun seems straightforward. Given a sentence such as (201):

201. He saw a snake in front of John.

one need only posit coreference, then attempt to find that reading in the sentence. This is essentially a semantic interpretation judgment, and is therefore qualitatively different from the task of judging the syntactic deviance of sentences whose interpretation is transparent:

202. ?That's the answer of which he got wind.

203. ?China is exciting for Sue and her friends to talk about visiting.

204. ?A picture will be published soon of the new prince.

Here, varying degrees of deviance make a yes/no judgment difficult, and this has been acknowledged in the extensive literature on the non-binary nature of grammaticality judgments (e.g., Bolinger, 1968;